arXiv:1411.2675v4 [math.OC] 18 Feb 2016

Process-Based Risk Measures and Risk-Averse Control of Observable and Partially Observable Discrete-Time Systems

Jingnan Fan* Andrzej Ruszczyński†

November 5, 2014; revised February 15,2016

Abstract

For controlled discrete-time stochastic processes we introduce a new class of dynamic risk measures, which we call process-based. Their main features are that they measure risk of processes that are functions of the history of a base process. We introduce a new concept of conditional stochastic time consistency and we derive the structure of process-based risk measures enjoying this property. We show that they can be equivalently represented by a collection of static law-invariant risk measures on the space of functions of the state of the base process. We apply this result to controlled Markov processes and we derive dynamic programming equations. Next, we consider partially observable Markov processes and we derive the structure of stochastically conditionally time-consistent risk measures in this case. We prove that the dynamic risk measures can be represented by a sequence of law invariant risk measures on the space of function of the observable part of the state. We also derive the corresponding dynamic programming equations.

Keywords: Dynamic Risk Measures; Time Consistency; Partially Observable Markov Processes; Dynamic Programming

1 Introduction

The main objective of this paper is to provide theoretical foundations of the theory of dynamic risk measures for controlled stochastic processes, in particular, Markov processes, with only partial observation of the state.

The theory of dynamic risk measures in discrete time has been intensively developed in the last 10 years (see [12, 42, 36, 38, 19, 10, 41, 2, 34, 25, 24, 11] and the references therein). The basic setting is the following: we have a probability space $(\Omega, \mathcal{F}, P)$, a filtration $\{\mathcal{F}_t\}_{t=1,\dots,T}$ with a trivial $\mathcal{F}_1$, and we define appropriate spaces $\mathcal{Z}_t$ of $\mathcal{F}_t$-measurable random variables, $t = 1, \dots, T$. For each $t = 1, \dots, T$, a mapping $\rho_{t,T} : \mathcal{Z}_T \to \mathcal{Z}_t$ is called a *conditional risk measure*. The central role in the theory is played by the concept of *time consistency*, which regulates relations between the mappings $\rho_{t,T}$ and $\rho_{s,T}$ for different s and t. One definition employed in

*Rutgers University, RUTCOR, Piscataway, NJ 08854, USA, Email: jf546@rutgers.edu

†Rutgers University, Department of Management Science and Information Systems, Piscataway, NJ 08854, USA, Email: rusz@rutgers.edu

the literature is the following: *for all* $Z, W \in \mathcal{Z}_T$, *if* $\rho_{t,T}(Z) \leq \rho_{t,T}(W)$ *then* $\rho_{s,T}(Z) \leq \rho_{s,T}(W)$ *for all* $s < t$. This can be used to derive recursive relations $\rho_{t,T}(Z) = \rho_t(\rho_{t+1,T}(Z))$, with simpler *one-step conditional risk mappings* $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$, $t = 1, \ldots, T-1$. Much effort has been devoted to derive dual representations of the conditional risk mapping and to study their evolution in various settings.

When applied to processes described by (controlled) kernels, in particular, to Markov processes, the theory of dynamic measures of risk encounters difficulties. The spaces $\mathcal{Z}_t$ are different for different t, and thus each one-step mapping ρ_t has different domain and range spaces. With $\mathcal{Z}_t$ containing all $\mathcal{F}_t$ measurable random variables, arbitrary dependence of ρ_t on the past is allowed. There is no easy way to explore stationarity of the process. Moreover, no satisfactory theory of law invariant dynamic risk measures exists, which would be suitable for Markov control problems (the restrictive definitions of law invariance employed in [26] and [43] lead to conclusions of limited practical usefulness, while the translation of the approach of [45] to the Markov case appears to be difficult). These difficulties are compounded in the case of controlled processes, when a control policy changes the probability measure on the space of paths of the process. Risk measurement of the entire family of processes defined by control policies is needed.

Motivated by these issues, in [40], we introduced a specific class of dynamic risk measures, which is well-suited for Markov problems. We postulated that the one-step conditional risk mappings ρ_t have a special form, which allows for their representation in terms of static risk measures on a space of functions defined on the state space of the Markov process. This restriction allowed for the development of dynamic programming equations and corresponding solution methods, which generalize the well-known results for expected value problems. Our ideas were successfully extended in [8, 7, 29, 44]. However, our construction of the Markov risk measures appeared somewhat arbitrary.

In this paper, we introduce and analyze a general class of risk measures, which we call *process-based*. We consider a controlled process $\{X_t\}_{t=1,\ldots,T}$ taking values in a Polish space $\mathcal{X}$ (the *state space*), whose conditional distributions are described by controlled history-dependent transition kernels $Q_t : \mathcal{X}^t \times \mathcal{U} \to \mathcal{P}(\mathcal{X})$, $t = 1, \ldots, T-1$, where $\mathcal{U}$ is a certain *control space*. Any history-dependent (measurable) control $u_t = \pi_t(x_1, \ldots, x_t)$ is allowed. In this setting, we are only interested in measuring risk of stochastic processes of the form $Z_t = c_t(X_t, u_t)$, $t = 1, \ldots, T$, where $c_t : \mathcal{X} \times \mathcal{U} \to \mathbb{R}$ can be any bounded measurable function. This restriction of the class of stochastic processes for which risk needs to be measured is one of the two cornerstones of our approach. The other cornerstone is our new concept of *stochastic conditional time-consistency*. It is more restrictive than the usual time consistency, because it involves conditional distributions and uses stochastic dominance rather than the pointwise order. These two foundations allow for the development of a theory of dynamic risk measures which can be fully described by a sequence of static law invariant risk measures on a space $\mathcal{V}$ of measurable functions on the state space $\mathcal{X}$. In the special case of controlled Markov processes, we derive the structure postulated in [40], thus providing its solid theoretical foundations. We also derive dynamic programming equations in a much more general setting than that of [40].

In the extant literature, three basic approaches to introduce risk aversion in Markov decision processes have been employed: utility functions (see, e.g., [22, 23, 15]), mean–variance models (see. e.g., [46, 18, 30, 1]), and entropic (exponential) models (see, e.g., [21, 31, 13, 16, 28, 5]). Our approach generalizes the utility and exponential models; the mean–variance models do not satisfy, in general, the monotonicity and time-consistency conditions, except the version of [9].

In the second part of the paper, we extend our theory to partially observable controlled Markov processes. In the expected-value case, this classical topic is covered in many mono-

graphs (see, e.g., [20, 6, 4] and the references therein). The standard approach is to consider the belief state space, involving the space of probability distributions of the unobserved part of the state. The recent article [17] provides the state-of-the-art setting. The risk averse case has been dealt, so far, with the use of the entropic risk measure [35, 5]. Our approach is based on representing a partially observable Markov process as a process with a history-dependent transition kernel. This allows for further specialization of our general theory. Our main result is that the conditional risk mappings can be equivalently modeled by static risk measures on the space of functions defined on the observable part of the state only. We also derive risk-averse dynamic programming equations for partially observable Markov models. In these equations, the state space comprises belief states and observable states, as in the expected value model, but the conditional expectation is repaced by a special transition risk mapping.

The paper is organized as follows. In sections 2.1–2.3, we formalize the basic model and review the concepts of risk measures and their time consistency. The first set of original results are presented in section 2.4; we introduce a new concept of stochastic conditional time consistency and we characterize the structure of dynamic risk measures enjoying this property (Theorem 2.10). In section 3, we extend these ideas to the case of controlled processes, and we prove Theorem 3.3 on the structure of measures of risk in this case. These results are further specialized to controlled Markov processes in section 4. We introduce the concept of a Markov risk measure and we derive its structure (Theorem 4.4). In section 4.2, we prove an analog of dynamic programming equations in this case. Section 5 is devoted to the analysis of the structure of risk measures for partially observable controlled Markov processes. The key result is Theorem 5.6, which simplifies modeling of conditional risk measures in this case. In section 5.5, we derive an analog of dynamic programming equations for risk-averse partially observable models. Finally, in section 6, we illustrate our approach on a machine replacement problem.

2 Risk Measures Based on Observable Processes

In this section we introduce fundamental concepts and properties of dynamic risk measures for uncontrolled stochastic processes in discrete time. In subsection 2.1 we set up our a probabilistic framework, and in subsections 2.2 and 2.3 we revisit some important concepts existing in the literature. In subsection 2.4, we introduce the new notion of stochastic conditional time consistency, which is a stronger requirement on the dynamic risk measure than the standard time consistency, and which is particularly useful for controlled stochastic processes. Based on this concept, we derive the structure of dynamic risk measures involving transition risk mappings: a family of static risk measures on the space of functions of a state.

2.1 Preliminaries

In all subsequent considerations, we work with a Borel subset $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ of a Polish space and the canonical measurable space $\left(\mathcal{X}^T, \mathcal{B}(\mathcal{X})^T\right)$ where T is a natural number and $\mathcal{B}(\mathcal{X})^T$ is the product σ-algebra. We use $\{X_t\}_{t=1,\dots,T}$ to denote the discrete-time process of canonical projections. We also define $\mathcal{H}_t = \mathcal{X}^t$ to be the space of possible histories up to time t, and we use h_t for a generic element of $\mathcal{H}_t$: a specific history up to time t. The random vector $(X_1, \cdots, X_t)$ will be denoted by H_t.

We assume that for all $t = 1, \dots, T-1$, the transition kernels, which describe the conditional

distribution of X_{t+1}, given $X_1, \cdots, X_t$, are measurable functions

$$Q_t : \mathcal{X}^t \to \mathcal{P}(\mathcal{X}), \quad t = 1, \ldots, T-1, \tag{1}$$

where $\mathcal{P}(\mathcal{X})$ is the set of probability measures on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$. These kernels, along with the initial distribution of X_1, define a unique probability measure P on the product space $\mathcal{X}^T$ with the product σ-algebra.

For a stochastic system described above, we consider a sequence of random variables $\{Z_t\}_{t=1,\ldots,T}$ taking values in $\mathbb{R}$; we assume that lower values of Z_t are preferred (e.g., Z_t represents a "cost" at time t). We require $\{Z_t\}_{t=1,\ldots,T}$ to be bounded and adapted to $\{\mathcal{F}_t\}_{t=1,\ldots,T}$ - the natural filtration generated by the process X. In order to facilitate our discussion, we introduce the following spaces:

$$\mathcal{Z}_t = \left\{ Z : \mathcal{X}^T \to \mathbb{R} \,\middle|\, Z \text{ is } \mathcal{F}_t\text{-measurable and bounded} \right\}, \quad t = 1, \ldots, T. \tag{2}$$

It is then equivalent to say that $Z_t \in \mathcal{Z}_t$. We also introduce the spaces

$$\mathcal{Z}_{t,T} = \mathcal{Z}_t \times \cdots \times \mathcal{Z}_T, \quad t = 1, \ldots, T.$$

Since Z_t is $\mathcal{F}_t$-measurable, a measurable function $\phi_t : \mathcal{X}^t \to \mathbb{R}$ exists such that $Z_t = \phi_t(X_1, \ldots, X_t)$. With a slight abuse of notation, we still use Z_t to denote this function.

2.2 Dynamic risk measures

In this subsection, we quickly review some definitions and concepts related to risk measures. Note that all relations (e.g., equality, inequality) between random variables are understood in the "everywhere" sense.

Definition 2.1. *A mapping $\rho_{t,T} : \mathcal{Z}_{t,T} \to \mathcal{Z}_t$, where $1 \le t \le T$, is called a* conditional risk measure*, if it satisfies the* monotonicity property*: for all $(Z_t, \ldots, Z_T)$ and $(W_t, \ldots, W_T)$ in $\mathcal{Z}_{t,T}$, if $Z_s \le W_s$, for all $s = t, \ldots, T$, then*

$$\rho_{t,T}(Z_t, \ldots, Z_T) \le \rho_{t,T}(W_t, \ldots, W_T).$$

Definition 2.2. *A conditional risk measure $\rho_{t,T} : \mathcal{Z}_{t,T} \to \mathcal{Z}_t$*
(i) *is* normalized *if $\rho_{t,T}(0, \ldots, 0) = 0$;*
(ii) *is* translation-invariant *if $\forall\, (Z_t, \ldots, Z_T) \in \mathcal{Z}_{t,T}$,*

$$\rho_{t,T}(Z_t, \ldots, Z_T) = Z_t + \rho_{t,T}(0, Z_{t+1}, \ldots, Z_T).$$

Throughout the paper, we assume all conditional risk measures to be at least normalized. Translation-invariance is a fundamental property, which will also be frequently used; under normalization, it implies that $\rho_{t,T}(Z_t, 0, \cdots, 0) = Z_t$.

Definition 2.3. *A conditional risk measure $\rho_{t,T}$ has the* local property *if for any $(Z_t, \ldots, Z_T) \in \mathcal{Z}_{t,T}$ and for any event $A \in \mathcal{F}_t$, we have*

$$\mathbb{1}_A \rho_{t,T}(Z_t, \ldots, Z_T) = \rho_{t,T}(\mathbb{1}_A Z_t, \ldots, \mathbb{1}_A Z_T).$$

The local property is important because it means that the conditional risk measure at time t restricted to any $\mathcal{F}_t$-event A is not influenced by the values that $Z_t, \ldots, Z_T$ take on A^c.

Definition 2.4. *A* dynamic risk measure *$\rho = \{\rho_{t,T}\}_{t=1,\ldots,T}$ is a sequence of conditional risk measures $\rho_{t,T} : \mathcal{Z}_{t,T} \to \mathcal{Z}_t$. We say that ρ is normalized, translation-invariant, or has the local property, if all $\rho_{t,T}$, $t = 1, \ldots, T$, satisfy the respective conditions of Definitions 2.2 or 2.3.*

2.3 Time Consistency

The notion of time consistency can be formulated in different ways, with weaker or stronger assumptions; but the key idea is that if one sequence of costs, compared to another sequence, has the same current cost and lower risk in the future, then it should have lower current risk. In this and the next subsection, we discuss two formulations of time consistency: the (now) standard one, and our new proposal specially suited for process-based measures. We also show how the tower property (the recursive relation between $\rho_{t,T}$ and $\rho_{t+1,T}$ that the time consistency implies) improves with the more refined time consistency concept. The following definition of time consistency was employed in [40].

Definition 2.5. *A dynamic risk measure* $\{\rho_{t,T}\}_{t=1,\dots,T}$ *is* time-consistent *if for any* $1 \le t < T$ *and for all* $(Z_t, \dots, Z_T), (W_t, \dots, W_T) \in \mathcal{Z}_t$*, the conditions*

$$\begin{cases} Z_t = W_t, \\ \rho_{t+1,T}(Z_{t+1}, \dots, Z_T) \le \rho_{t+1,T}(W_{t+1}, \dots, W_T), \end{cases}$$

imply that $\rho_{t,T}(Z_t, \dots, Z_T) \le \rho_{t,T}(W_t, \dots, W_T)$.

It turns out that a translation-invariant and time-consistent dynamic risk measure can be decomposed into and then reconstructed from so-called *one-step conditional risk mappings*.

Theorem 2.6 ([40])**.** *A dynamic risk measure* $\{\rho_{t,T}\}_{t=1,\dots,T}$ *is translation-invariant and time-consistent if and only if there exist mappings* $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t, t = 1, \dots, T-1$*, satisfying the monotonicity and normalization properties, called* one-step conditional risk mappings*, such that for all* $t = 1, \dots, T-1$*,*

$$\rho_{t,T}(Z_t, \dots, Z_T) = Z_t + \rho_t(\rho_{t+1,T}(Z_{t+1}, \dots, Z_T)). \tag{3}$$

In general, time consistency does not imply the local property, unless additional conditions are satisfied.

Conceptually, the one-step conditional risk mappings play a similar role to one-step conditional expectations, and will be very useful when an analog of the tower property is involved. At this stage, without further refinement of assumptions, it remains a fairly abstract and general object that is hard to characterize. In [40], a special form of this one-step conditional risk mappings was imposed, which was well suited for Markovian applications, but it is unclear whether other forms of such mappings exist. In order to gain deeper understanding of these concepts, in the following we introduce a stronger notion of time consistency, and we argue that any one-step conditional risk mapping is of the form postulated in [40]. To this end, we use the particular structure of the space $\left(\mathcal{X}^T, \mathcal{B}(\mathcal{X})^T\right)$ and the way a probability measure is defined on this space.

2.4 Stochastic Conditional Time-Consistency and Transition Risk Mappings

We now refine the concept of time-consistency for process-based risk measures.

Definition 2.7. *A dynamic risk measure* $\{\rho_{t,T}\}_{t=1,\dots,T}$ *is* stochastically conditionally time-consistent with respect to $\{Q_t\}_{t=1,\dots,T-1}$ *if for any* $1 \le t \le T-1$*, for any* $h_t \in \mathcal{X}^t$*, and for all* $(Z_t, \dots, Z_T)$, $(W_t, \dots, W_T) \in \mathcal{Z}_{t,T}$*, the conditions*

$$\begin{cases} Z_t(h_t) = W_t(h_t), \\ (\rho_{t+1,T}(Z_{t+1}, \dots, Z_T) \mid H_t = h_t) \preceq_{\text{st}} (\rho_{t+1,T}(W_{t+1}, \dots, W_T) \mid H_t = h_t), \end{cases} \tag{4}$$

imply

$$\rho_{t,T}(Z_t, \dots, Z_T)(h_t) \le \rho_{t,T}(W_t, \dots, W_T)(h_t), \tag{5}$$

where the relation $\preceq_{\mathrm{st}}$ *stands for the conditional stochastic order understood as follows:*

$$Q_t(h_t)\Big(\{\, x \,|\, \rho_{t+1,T}(Z_{t+1}, \dots, Z_T)(h_t, x) > \eta \,\}\Big) \\ \le Q_t(h_t)\Big(\{\, x \,|\, \rho_{t+1,T}(W_{t+1}, \dots, W_T)(h_t, x) > \eta \,\}\Big), \quad \forall\, \eta \in \mathbb{R}.$$

When the choice of the underlying transition kernels is clear from the context, we will simply say that the dynamic risk measure is stochastically conditionally time-consistent.

Proposition 2.8. *The stochastic conditional time-consistency implies the time-consistency and the local property.*

Proof. If $\{\rho_{t,T}\}_{t=1,\dots,T}$ is stochastically conditionally time-consistent, then it satisfies Definition 2.5 and is time-consistent.

Let us prove by induction on t from T down to 1 that $\rho_{t,T}$ satisfy the local property. Clearly, $\rho_{T,T}$ does: if $A \in \mathcal{F}_T$, then $[\mathbb{1}_A \rho_{T,T}](Z_T) = \mathbb{1}_A Z_T = \rho_{T,T}(\mathbb{1}_A Z_T)$.

Suppose $\rho_{t+1,T}$ satisfies the local property for some $1 \le t < T$, and consider any $A \in \mathcal{F}_t$, any $h_t \in \mathcal{X}^t$, and any $(Z_t, \dots, Z_T) \in \mathcal{Z}_{t,T}$. Two cases may occur.

- If $\mathbb{1}_A(h_t) = 0$, then $[\mathbb{1}_A Z_t](h_t) = 0$ and the local property for $t+1$ yields:

  $$[\rho_{t+1,T}(\mathbb{1}_A Z_{t+1}, \dots, \mathbb{1}_A Z_T)](h_t, \cdot) = [\mathbb{1}_A \rho_{t+1,T}(Z_{t+1}, \dots, Z_T)](h_t, \cdot) = 0.$$

  By stochastic conditional time consistency,

  $$\rho_{t,T}(\mathbb{1}_A Z_t, \mathbb{1}_A Z_{t+1}, \dots, \mathbb{1}_A Z_T)(h_t) = \rho_{t,T}(0, \dots, 0)(h_t) = 0.$$

- If $\mathbb{1}_A(h_t) = 1$, then $[\mathbb{1}_A Z_t](h_t) = Z_t(h_t)$ and the local property for $t+1$ implies that

  $$[\rho_{t+1,T}(\mathbb{1}_A Z_{t+1}, \dots, \mathbb{1}_A Z_T)](h_t, \cdot) = [\mathbb{1}_A \rho_{t+1,T}(Z_{t+1}, \dots, Z_T)](h_t, \cdot) \\ = \rho_{t+1,T}(Z_{t+1}, \dots, Z_T)(h_t, \cdot).$$

  By stochastic conditional time consistency, $\rho_{t,T}(\mathbb{1}_A Z_t, \dots, \mathbb{1}_A Z_T)(h_t) = \rho_{t,T}(Z_t, \dots, Z_T)(h_t)$.

In both cases, $\rho_{t,T}(\mathbb{1}_A Z_t, \dots, \mathbb{1}_A Z_T)(h_t) = [\mathbb{1}_A \rho_{t,T}(Z_t, \dots, Z_T)](h_t)$. □

Generally, the definition of dynamic risk measures in Section 2.2 and the definition of time consistency property in Section 2.3 are valid with any filtration on an underlying probability space $(\Omega, \mathcal{F}, P)$, instead of process-generated filtration. However, from now on, we only consider dynamic risk measures defined with a filtration generated by a *specific* process, because we are interested in those risk measures which can be evaluated on each specific history path. That is why we call these risk measures "process-based."

The following proposition shows that the stochastic conditional time consistency implies that the one-step risk mappings ρ_t can be equivalently represented by static law-invariant risk measures on $\mathcal{V}$, the set of all bounded measurable functions on $\mathcal{X}$. We first slightly refine the standard concept of law invariance.

Definition 2.9. *A risk measure $r : \mathcal{V} \to \mathbb{R}$ is law invariant with respect to the probability measure q on $(\mathcal{V}, \mathcal{B}(\mathcal{V}))$, if for all $V, W \in \mathcal{V}$*

$$V \overset{q}{\sim} W \Rightarrow r(V) = r(W),$$

where $V \overset{q}{\sim} W$ means that $q\{V \leq \eta\} = q\{W \leq \eta\}$ for all $\eta \in \mathbb{R}$.

We can now state the main result of this section.

Theorem 2.10. *A process-based dynamic risk measure $\{\rho_{t,T}\}_{t=1,\dots,T}$ is translation-invariant and stochastically conditionally time-consistent if and only if functionals $\sigma_t : \mathrm{graph}(Q_t) \times \mathcal{V} \to \mathbb{R}$, $t = 1, \dots, T-1$, exist, such that*
(i) *for all $t = 1, \dots, T-1$ and all $h_t \in \mathcal{X}^t$, the functional $\sigma_t(h_t, Q_t(h_t), \cdot)$ is a normalized, monotonic, and law-invariant risk measure on $\mathcal{V}$ with respect to the distribution $Q_t(h_t)$;*
(ii) *for all $t = 1, \dots, T-1$, for all $(Z_t, \dots, Z_T) \in \mathcal{Z}_{t,T}$, and for all $h_t \in \mathcal{X}^t$,*

$$\rho_{t,T}(Z_t, \dots, Z_T)(h_t) = Z_t(h_t) + \sigma_t(h_t, Q_t(h_t), \rho_{t+1,T}(Z_{t+1}, \dots, Z_T)(h_t, \cdot)). \tag{6}$$

Moreover, for all $t = 1, \dots, T-1$, σ_t is uniquely determined by $\rho_{t,T}$ as follows: for every $h_t \in \mathcal{X}^t$ and every $v \in \mathcal{V}$,

$$\sigma_t(h_t, Q_t(h_t), v) = \rho_{t,T}(0, V, 0, \dots, 0)(h_t), \tag{7}$$

where $V \in \mathcal{Z}_{t+1}$ satisfies the equation $V(h_t, \cdot) = v(\cdot)$, and can be arbitrary elsewhere.

Proof. Assume $\{\rho_{t,T}\}_{t=1,\dots,T}$ is translation-invariant and stochastically conditionally time-consistent. For any $h_t \in \mathcal{X}^t$, the formula (7) defines a normalized and monotonic risk measure on the space $\mathcal{V}$. Setting

$$v(x) = \rho_{t+1,T}(Z_{t+1}, \dots, Z_T)(h_t, x), \quad \forall\, x \in \mathcal{X},$$
$$V(h_{t+1}) = \begin{cases} v(x), & \text{if } h_{t+1} = (h_t, x), \\ 0, & \text{otherwise,} \end{cases}$$

we obtain, by translation invariance and normalization,

$$\rho_{t+1,T}(V, 0, \dots, 0)(h_t, \cdot) = V(h_t, \cdot) = \rho_{t+1,T}(Z_{t+1}, \dots, Z_T)(h_t, \cdot).$$

Thus, by the translation property and stochastic conditional time consistency,

$$\begin{aligned} \rho_{t,T}(Z_t, \dots, Z_T)(h_t) &= Z_t(h_t) + \rho_{t,T}(0, Z_{t+1}, \dots, Z_T)(h_t) \\ &= Z_t(h_t) + \rho_{t,T}(0, V, 0, \dots, 0)(h_t) \\ &= Z_t(h_t) + \sigma_t(h_t, v). \end{aligned}$$

This chain of relations proves also the uniqueness of σ_t. We need only verify the postulated law invariance of $\sigma_t(h_t, Q_t(h_t), \cdot)$. If $V, V' \in \mathcal{Z}_{t+1}$ have the same conditional distribution, given h_t, then Definition 2.7 implies that $\rho_{t,T}(0, V, 0, \dots, 0)(h_t) = \rho_{t,T}(0, V', 0, \dots, 0)(h_t)$, and law invariance follows from (7).

On the other hand, if such transition risk mappings exist, then $\{\rho_{t,T}\}_{t=1,\dots,T}$ is stochastically conditionally time-consistent by the monotonicity and law invariance of $\sigma(h_t, \cdot)$. We can now use (6) to obtain for any $t = 1, \dots, T-1$, and for all $h_t \in \mathcal{X}^t$ the following identity:

$$\begin{aligned} \rho_{t,T}(0, Z_{t+1}, \dots, Z_T)(h_t) &= \sigma_t(h_t, \rho_{t+1,T}(Z_{t+1}, \dots, Z_T)(h_t, \cdot)) \\ &= \rho_{t,T}(Z_t, \dots, Z_T)(h_t) - Z_t(h_t), \end{aligned}$$

which is translation invariance of $\rho_{t,T}$. □

Remark 2.11. *With a slight abuse of notation, we included the distribution $Q_t(h_t)$ as an argument of the transition risk mapping in view of the application to controlled processes.*

Example 2.12. *In the theory of risk-sensitive Markov decision processes, the following family of* entropic risk measures *is employed (see [21, 35, 31, 14, 39, 13, 16, 28, 5]):*

$$\rho_{t,T}(Z_t, \dots, Z_T) = \frac{1}{\gamma} \ln\Big(\mathbb{E}\Big[\exp\Big(\gamma \textstyle\sum_{s=t}^{T} Z_s\Big) \,\Big|\, \mathcal{F}_t\Big]\Big), \quad t = 1, \dots, T, \quad \gamma > 0.$$

It is stochastically conditionally time-consistent, and corresponds to the transition risk mapping

$$\sigma_t(h_t, q, v) = \frac{1}{\gamma} \ln\left(\mathbb{E}_q[e^{\gamma v}]\right) = \frac{1}{\gamma} \ln\Big(\int_{\mathcal{X}} e^{\gamma v(x)}\, q(dx)\Big), \quad \gamma > 0. \tag{8}$$

We could also make γ in (8) *dependent on the time t, the current state x_t, or even the entire history h_t, and still obtain a stochastically conditionally time-consistent dynamic risk measure.*

Example 2.13. *The following transition risk mapping satisfies the condition of Theorem 2.10 and corresponds to a stochastically conditionally time-consistent dynamic risk measure:*

$$\sigma_t(h_t, q, v) = \int_{\mathcal{X}} v(s)\, q(ds) + \varkappa_t(h_t) \left(\int_{\mathcal{X}} \Big[\Big(v(s) - \int_{\mathcal{X}} v(s')\, q(ds')\Big)_+\Big]^p\, q(ds)\right)^{1/p}, \tag{9}$$

where $\varkappa_t : \mathcal{X}^t \to [0, 1]$ is a measurable function, and $p \in [1, +\infty)$. It is an analogue of the static mean–semideviation measure of risk, whose consistency with stochastic dominance is well–known [32, 33]. In the construction of a dynamic risk measure, we use $q = Q_t(h_t)$. If $\varkappa_t$ depends on x_t only, the mapping (9) *corresponds to a Markov risk measure discussed in sections 4.1 and 4.2.*

Example 2.14. *Our use of the stochastic dominance relation in the definition of stochastic conditional time consistency rules out some candidates for transition risk mappings. Suppose $\sigma_t(h_t, q, v) = v(x_1)$, where $x_1 \in \mathcal{X}$ is a selected state. Such a mapping is a coherent measure of risk, as a function of the last argument, and may be law invariant. In particular, it is law invariant with $\mathcal{X} = \{x_1, x_2\}$, $q(x_1) = 1/3$, $q(x_2) = 2/3$, $v(x_1) = 3$, $v(x_2) = 1$, $w(x_1) = 2$, $w(x_2) = 4$. For this mapping, we have $v \preceq_{\mathrm{st}} w$ under q, but $\sigma_t(h_t, q, v) > \sigma_t(h_t, q, w)$, and thus the condition of stochastic conditional time consistency is violated. This is due to the fact that the probability of reaching x_1, no matter how small, does not affect the value of the risk measure. We consciously exclude such cases, because in controlled systems, to be discussed in the next section, the second argument (q) is the only one that depends on our decisions. It should be included in the definition of our preferences, if practically meaningful results are to be obtained. have to*

3 Risk Measures for Controlled Stochastic Processes

We now extend the setting of Section 2 by allowing the kernels (1) to depend on control variables u_t.

3.1 The Model

We still work with the process $\{X_t\}_{t=1,\dots,T}$ on the space $\mathcal{X}^T$ and introduce a Polish control space $(\mathcal{U}, \mathcal{B}(\mathcal{U}))$. At each time t, we observe the state x_t and then apply a control $u_t \in \mathcal{U}$. We assume

that the admissible control sets and the transition kernels (conditional distributions of the next state) depend on all currently-known state and control values. More precisely, we make the following assumptions:

1. For all $t = 1, \dots, T$, we require that $u_t \in \mathcal{U}_t(x_1, u_1, \dots, x_{t-1}, u_{t-1}, x_t)$, where $\mathcal{U}_t : \mathcal{G}_t \rightrightarrows \mathcal{U}$ is a measurable multifunction, and $\mathcal{G}_1, \dots, \mathcal{G}_T$ are the sets of histories of all currently-known state and control values before applying each control:

$$\begin{cases} \mathcal{G}_1 = \mathcal{X}, \\ \mathcal{G}_{t+1} = \operatorname{graph}(\mathcal{U}_t) \times \mathcal{X} \subseteq (\mathcal{X} \times \mathcal{U})^t \times \mathcal{X}, \quad t = 1, \dots, T-1; \end{cases}$$

2. For all $t = 1, \dots, T$, the control-dependent transition kernels

$$Q_t : \operatorname{graph}(\mathcal{U}_t) \to \mathcal{P}(\mathcal{X}), \quad t = 1, \dots, T-1, \tag{10}$$

are measurable, and for all $t = 1, \dots, T-1$, for all $(x_1, u_1, \dots, x_t, u_t) \in \operatorname{graph}(\mathcal{U}_t)$, $Q_t(x_1, u_1, \dots, x_t, u_t)$ describes the conditional distribution of X_{t+1}, given currently-known states and controls.

The reason why we allow $\mathcal{U}_t$ and Q_t to depend on all known state and control values is that in partially observable Markov decision problems, which we study in Section 5, the conditional distributions of the next observable part of the state does depend on all past controls. We plan to apply the results of the present section to the partially observable case.

For this controlled process, a (deterministic) *history-dependent admissible policy* $\pi = (\pi_1, \dots, \pi_T)$ is a sequence of measurable selectors, called *decision rules*, $\pi_t : \mathcal{G}_t \to \mathcal{U}$ such that $\pi_t(g_t) \in \mathcal{U}_t(g_t)$ for all $g_t \in \mathcal{G}_t$. We can easily prove by induction on t that for such an admissible policy π, each π_t reduces to a measurable function on $\mathcal{X}^t$, as $u_s = \pi_s(h_s)$ for all $s = 1, \dots, t-1$. We are still using π_s to denote the decision rule, although it is a different function, formally; it will not lead to any misunderstanding. The set of admissible policies is

$$\begin{aligned} \Pi := \{ \pi = (\pi_1, \dots, \pi_T) \mid \\ \forall t,\ \pi_t(x_1, \dots, x_t) \in \mathcal{U}_t(x_1, \pi_1(x_1), \dots, x_{t-1}, \pi_{t-1}(x_1, \dots, x_{t-1}), x_t) \}. \end{aligned} \tag{11}$$

For any fixed policy $\pi \in \Pi$, the transition kernels can be rewritten as measurable functions from $\mathcal{X}^t$ to $\mathcal{P}(\mathcal{X})$:

$$Q_t^\pi : (x_1, \dots, x_t) \mapsto Q_t(x_1, \pi_1(x_1), \dots, x_t, \pi_t(x_1, \dots, x_t)),\ t = 1, \dots, T-1, \tag{12}$$

just like the transition kernels of the uncontrolled case given in (1), but indexed by π. Thus for any policy $\pi \in \Pi$, we can consider $\{X_t\}_{t=1,\dots,T}$ as an "uncontrolled" process, on the probability space $(\mathcal{X}^T, \mathcal{B}(\mathcal{X})^T, P^\pi)$ with P^π defined by $\{Q_t^\pi\}_{t=1,\dots,T-1}$, which is adapted to the policy-independent filtration $\{\mathcal{F}_t\}_{t=1,\dots,T}$. As before and throughout this paper, $h_t \in \mathcal{X}^t$ will stand for $(x_1, \dots, x_t)$.

We still use the same spaces $\mathcal{Z}_t$, $t = 1, \dots, T$ as defined in (2) for the costs incurred at each stage; these spaces also allow us to consider control-dependent costs as collections of policy-indexed costs in $\mathcal{Z}_{1,T}$. Thus, we are able to define and analyze (time-consistent) dynamic risk measures ρ^π for each fixed $\pi \in \Pi$, as in Section 2. Note that ρ^π are defined on the same spaces independently of π, because the filtration and the spaces $\mathcal{Z}_t$, $t = 1, \dots, T$, are not dependent on π; however, we do need to index the measures of risk by the policy π, because the transition kernels and, consequently, the probability measure on the space $\mathcal{X}^T$, depend on π.

3.2 Stochastic Conditional Time Consistency and Transition Risk Mappings

We need to compare risk levels among different policies, so a meaningful order among the risk measures (ρ^π), $\pi \in \Pi$ is needed. It turns out that our concept of stochastic conditional time-consistency can be extended to this setting.

Definition 3.1. *A family of process-based dynamic risk measures* $\{\rho^\pi_{t,T}\}^{\pi\in\Pi}_{t=1,\dots,T-1}$ *is* stochastically conditionally time-consistent *if for any* $\pi, \pi' \in \Pi$*, for any* $1 \le t < T$*, for all* $h_t \in \mathcal{X}^t$*, all* $(Z_t, \dots, Z_T) \in \mathcal{Z}_{t,T}$ *and all* $(W_t, \dots, W_T) \in \mathcal{Z}_{t,T}$*, the conditions*

$$\begin{cases} Z_t(h_t) = W_t(h_t), \\ (\rho^\pi_{t+1,T}(Z_{t+1}, \dots, Z_T) \mid H^\pi_t = h_t) \preceq_{\mathrm{st}} (\rho^{\pi'}_{t+1,T}(W_{t+1}, \dots, W_T) \mid H^{\pi'}_t = h_t), \end{cases}$$

imply

$$\rho^\pi_{t,T}(Z_t, \dots, Z_T)(h_t) \le \rho^{\pi'}_{t,T}(W_t, \dots, W_T)(h_t).$$

Remark 3.2. *As in Definition 2.7, the conditional stochastic order "*$\preceq_{\mathrm{st}}$*" should be understood as follows: for all* $\eta \in \mathbb{R}$ *we have*

$$Q^\pi_t(h_t)\Big(\{\, x \,|\, \rho^\pi_{t+1,T}(Z_{t+1}, \dots, Z_T)(h_t, x) > \eta \,\}\Big) \\ \le Q^{\pi'}_t(h_t)\Big(\{\, x \,|\, \rho^{\pi'}_{t+1,T}(W_{t+1}, \dots, W_T)(h_t, x) > \eta \,\}\Big).$$

The above definition plays a pivotal role in building an intrinsic connection among dynamic risk measures $(\rho^\pi)_{\pi\in\Pi}$, as we explain it below. Before passing to the details, we can say in short that the same transition risk mappings as in the uncontrolled case are the only possible structures of such risk measures.

If a family of process-based dynamic risk measures $\{\rho^\pi_{t,T}\}^{\pi\in\Pi}_{t=1,\dots,T-1}$ is stochastically conditionally time-consistent, then for each fixed $\pi \in \Pi$ the process-based dynamic risk measure $\{\rho^\pi_{t,T}\}_{t=1,\dots,T-1}$ is stochastically conditionally time-consistent, as defined in Definition 2.7. By virtue of Proposition 2.10, for each $\pi \in \Pi$, there exist functionals $\sigma^\pi_t : \mathrm{graph}(Q^\pi_t) \times \mathcal{V} \to \mathbb{R}$, $t = 1 \dots T-1$, such that for all $t = 1, \dots, T-1$, all $h_t \in \mathcal{X}^t$, the functional $\sigma^\pi_t(h_t, Q^\pi_t(h_t), \cdot)$ is a law-invariant risk measure on $\mathcal{V}$ with respect to the distribution $Q^\pi_t(h_t)$ and

$$\rho^\pi_{t,T}(Z_t, \dots, Z_T)(h_t) = Z_t(h_t) + \sigma^\pi_t(h_t, Q^\pi_t(h_t), \rho^\pi_{t+1,T}(Z_{t+1}, \dots, Z_T)(h_t, \cdot)), \quad \forall h_t \in \mathcal{X}^t.$$

Consider any $\pi, \pi' \in \Pi$, $h_t \in \mathcal{X}^t$, and $(Z_t, \dots, Z_T) \in \mathcal{Z}_{t,T}$, $(W_t, \dots, W_T) \in \mathcal{Z}_{t,T}$ such that

$$\begin{cases} Z_t(h_t) = W_t(h_t), \\ Q^\pi_t(h_t) = Q^{\pi'}_t(h_t), \\ \rho^\pi_{t+1,T}(Z_{t+1}, \dots, Z_T)(h_t, \cdot) = \rho^{\pi'}_{t+1,T}(W_{t+1}, \dots, W_T)(h_t, \cdot). \end{cases}$$

Then we have

$$(\rho^\pi_{t+1,T}(Z_{t+1}, \dots, Z_T) \mid H^\pi_t = h_t) \sim_{\mathrm{st}} (\rho^{\pi'}_{t+1,T}(W_{t+1}, \dots, W_T) \mid H^{\pi'}_t = h_t),$$

where the relation $\sim_{\mathrm{st}}$ means that both $\preceq_{\mathrm{st}}$ and $\succeq_{\mathrm{st}}$ are true; in other words, equality in law. Because of the stochastic conditional time-consistency,

$$\rho^\pi_{t,T}(Z_t, \dots, Z_T)(h_t) = \rho^{\pi'}_{t,T}(W_t, \dots, W_T)(h_t),$$

whence

$$\sigma_t^{\pi}(h_t, Q_t^{\pi}(h_t), \rho_{t+1,T}^{\pi}(Z_{t+1}, \dots, Z_T)(h_t, \cdot)) \\ = \sigma_t^{\pi'}(h_t, Q_t^{\pi'}(h_t), \rho_{t+1,T}^{\pi'}(W_{t+1}, \dots, W_T)(h_t, \cdot)).$$

This proves that in fact σ^{π} does not depend on π, and all dependence on π is carried by the controlled kernel Q_t^{π}. This is a highly desirable property, when we apply dynamic risk measures to a control problem. We summarize this important observation in the following theorem, which extends Theorem 2.10 to the case of controlled processes.

Theorem 3.3. *A family of process-based dynamic risk measures $\{\rho_{t,T}^{\pi}\}_{t=1,\dots,T}^{\pi \in \Pi}$ is translation-invariant and stochastically conditionally time-consistent if and only if there exist functionals*

$$\sigma_t : \left\{ \bigcup_{\pi \in \Pi} \operatorname{graph}(Q_t^{\pi}) \right\} \times \mathcal{V} \to \mathbb{R}, \quad t = 1 \dots T-1,$$

such that
(i) *For all $t = 1, \dots, T-1$ and all $h_t \in \mathcal{X}^t$, $\sigma_t(h_t, \cdot, \cdot)$ is normalized and has the following property of strong monotonicity with respect to stochastic dominance:*

$$\forall q^1, q^2 \in \{ Q_t^{\pi}(h_t) : \pi \in \Pi \}, \ \forall v^1, v^2 \in \mathcal{V}, \\ (v^1; q^1) \preceq_{st} (v^2; q^2) \implies \sigma_t(h_t, q^1, v^1) \le \sigma_t(h_t, q^2, v^2),$$

where $(v; q) = q \circ v^{-1}$ means "the distribution of v under q;"
(ii) *For all $\pi \in \Pi$, for all $t = 1, \dots, T-1$, for all $(Z_t, \dots, Z_T) \in \mathcal{Z}_{t,T}$, and for all $h_t \in \mathcal{X}^t$,*

$$\rho_{t,T}^{\pi}(Z_t, \dots, Z_T)(h_t) = Z_t(h_t) + \sigma_t(h_t, Q_t^{\pi}(h_t), \rho_{t+1,T}^{\pi}(Z_{t+1}, \dots, Z_T)(h_t, \cdot)). \tag{13}$$

Moreover, for all $t = 1, \dots, T-1$, σ_t is uniquely determined by $\rho_{t,T}$ as follows: for every $h_t \in \mathcal{X}^t$, for every $q \in \{ Q_t^{\pi}(h_t) : \pi \in \Pi \}$, and for every $v \in \mathcal{V}$,

$$\sigma_t(h_t, q, v) = \rho_{t,T}^{\pi}(0, V, 0, \dots, 0)(h_t), \tag{14}$$

where π is any admissible policy such that $q = Q_t^{\pi}(h_t)$, and $V \in \mathcal{Z}_{t+1}$ satisfies the equation $V(h_t, \cdot) = v(\cdot)$, and can be arbitrary elsewhere.

Proof. We have shown the existence of $\{\sigma_t\}_{t=1,\dots,T}$ satisfying (13) and (14) in the discussion preceding the theorem. We can verify the strong law-invariance by (14) and Definition 3.1. □

It follows that the transition risk mappings of Examples 2.12 and 2.13 are perfectly suitable transition risk mappings for controlled processes as well.

4 Application to Controlled Markov Systems

Our results can be further specialized to the case when $\{X_t\}$ is a controlled Markov system, in which we assume the following conditions:
(i) The admissible control sets are measurable multifunctions of the current state, i.e., $\mathcal{U}_t : \mathcal{X} \rightrightarrows \mathcal{U}, \quad t = 1, \dots, T$;
(ii) The dependence in the transition kernel (10) on the history is carried only through the last state and control: $Q_t : \operatorname{graph}(\mathcal{U}_t) \to \mathcal{P}(\mathcal{X})$, $t = 1, \dots, T-1$;

(iii) The step-wise costs are only dependent on the current state and control: $Z_t = c_t(x_t, u_t)$, $t = 1, \dots, T$, where $c_t : \text{graph}(\mathcal{U}_t) \to \mathbb{R}$, $t = 1, \dots, T$ are measurable bounded functions.

Let Π be the set of admissible *history-dependent policies*:

$$\Pi := \{\, \pi = (\pi_1, \dots, \pi_T) \,|\, \forall t,\ \pi_t(x_1, \dots, x_t) \in \mathcal{U}_t(x_t) \,\}.$$

To alleviate notation, for all $\pi \in \Pi$ and for all measurable $c = (c_1, \dots, c_T)$, we write

$$v_t^{c,\pi}(h_t) := \rho_{t,T}^{\pi}(c_t(X_t, \pi_t(H_t)), \dots, c_T(X_T, \pi_T(H_T)))(h_t).$$

The following result is the direct consequence of Theorem 3.3 in the Markovian case.

Corollary 4.1. *For a controlled Markov system, a family of process-based dynamic risk measures $\{\rho_{t,T}^{\pi}\}_{t=1,\dots,T}^{\pi\in\Pi}$ is translation-invariant and stochastically conditionally time-consistent if and only if functionals*

$$\sigma_t : \Big\{(h_t, Q_t(x_t, u)) : h_t \in \mathcal{X}^t, u \in \mathcal{U}_t(x_t)\Big\} \times \mathcal{V} \to \mathbb{R}, \quad t = 1 \dots T-1,$$

exist, such that
(i) *For all $t = 1, \dots, T-1$ and all $h_t \in \mathcal{X}^t$, $\sigma_t(h_t, \cdot, \cdot)$ is normalized and strongly monotonic with respect to stochastic dominance on $\{\, Q_t(x_t, u) : u \in \mathcal{U}_t(x_t) \,\}$;*
(ii) *For all $\pi \in \Pi$, for all bounded measurable c, for all $t = 1, \dots, T-1$, and for all $h_t \in \mathcal{X}^t$,*

$$v_t^{c,\pi}(h_t) = c_t(x_t, \pi_t(h_t)) + \sigma_t\Big(h_t, Q_t(x_t, \pi_t(h_t)), v_{t+1}^{c,\pi}(h_t, \cdot)\Big). \tag{15}$$

Proof. To verify the "if and only if" statement, we can show that (14) is true if σ_t satisfies (15) for all measurable bounded c. □

4.1 Markov Risk Measures

Consider a Markov policy π composed of state-dependent measurable decision rules $\pi_t : \mathcal{X} \mapsto \mathcal{U}$, $t = 1, \dots, T$. Because of the Markov property of the transition kernels, for a Markov policy π, the future evolution of the process $\{X_\tau\}_{\tau=t,\dots,T}$ is solely dependent on the current state x_t, so is the distribution of the future costs $c_\tau(X_\tau, \pi_\tau(X_\tau))$, $\tau = t, \dots, T$. Therefore, it is reasonable to assume that the dependence of the conditional risk measure on the history is also carried by the current state only.

Definition 4.2. *A family of process-based dynamic risk measures $\{\rho_{t,T}^{\pi}\}_{t=1,\dots,T}^{\pi\in\Pi}$ for a controlled Markov system is* Markov *if for all Markov policies $\pi \in \Pi$, for all measurable $c = (c_1, \dots, c_T)$, and for all $h_t = (x_1, \dots, x_t)$ and $h_t' = (x_1', \dots, x_t')$ in $\mathcal{X}^t$ such that $x_t = x_t'$, we have $v_t^{c,\pi}(h_t) = v_t^{c,\pi}(h_t')$.*

Proposition 4.3. *Under translation invariance and stochastic conditional time consistency, $\{\rho_{t,T}^{\pi}\}_{t=1,\dots,T}^{\pi\in\Pi}$ is Markov if and only if the dependence of σ_t on h_t is carried only by x_t, for all $t = 1, \dots, T-1$.*

Proof. The Markov property in Definition 4.2 implies that for all $t = 1, \dots, T-1$, for all $h_t, h_t' \in \mathcal{X}^t$ such that $x_t = x_t'$, for all $u \in \mathcal{U}_t(x_t)$ and for all $v \in \mathcal{V}$, there exists a Markov $\pi \in \Pi$ such that $\pi_t(x_t) = u$. By setting $c = (0, \dots, 0, c_{t+1}, 0, \dots, 0)$ with $c_{t+1} : (x', u') \mapsto v(x')$, we obtain

$$\sigma_t(h_t, Q_t(x_t, u), v) = v_t^{c,\pi}(h_t) = v_t^{c,\pi}(h_t') = \sigma_t(h_t', Q_t(x_t, u), v).$$

Therefore, σ_t is indeed memoryless, that is, its dependence on h_t is carried by x_t only.

If σ_t, $t = 1, \dots, T-1$ are all memoryless, we can prove by induction backward in time that for all $t = T, \dots, 1$, $v_t^{c,\pi}(h_t) = v_t^{c,\pi}(h_t')$ for all Markov π and all $h_t, h_t' \in \mathcal{X}^t$ such that $x_t = x_t'$. □

Theorem 4.4. *For a controlled Markov system, a family of process-based dynamic risk measures $\{\rho_{t,T}^{\pi}\}_{t=1,\dots,T}^{\pi\in\Pi}$ is translation-invariant, stochastically conditionally time-consistent, and Markov, if and only if there exist functionals*

$$\sigma_t : \{(x, Q_t(x,u)) : x \in \mathcal{X}, u \in \mathcal{U}_t(x)\} \times \mathcal{V} \to \mathbb{R}, \quad t = 1 \dots T-1,$$

where $\mathcal{V}$ is the set of bounded measurable functions on $\mathcal{X}$, such that:
(i) *For all $t = 1, \dots, T-1$ and all $x \in \mathcal{X}$, $\sigma_t(x, \cdot, \cdot)$ is normalized and strongly monotonic with respect to stochastic dominance on $\{ Q_t(x,u) : u \in \mathcal{U}_t(x) \}$;*
(ii) *For all $\pi \in \Pi$, for all measurable bounded c, for all $t = 1, \dots, T-1$, and for all $h_t \in \mathcal{X}^t$,*

$$v_t^{c,\pi}(h_t) = c_t(x_t, \pi_t(h_t)) + \sigma_t\Big(x_t, Q_t(x_t, \pi_t(h_t)), v_{t+1}^{c,\pi}(h_t, \cdot)\Big). \tag{16}$$

Theorem 4.4 provides us with a simple recursive formula (16) for the evaluation of risk of a Markov policy π:

$$\begin{aligned} &v_T^{c,\pi}(x) = c_T(x, \pi_T(x)), \quad x \in \mathcal{X}, \\ &v_t^{c,\pi}(x) = c_t(x, \pi_t(x)) + \sigma_t(x, Q_t(x, \pi_t(x)), v_{t+1}^{c,\pi}), \quad x \in \mathcal{X}, \quad t = T-1, \dots, 1. \end{aligned}$$

It involves calculation of the values of functions $v_t^{c,\pi}(\cdot)$ on the state space $\mathcal{X}$.

4.2 Dynamic Programming

In this section, we fix the cost functions $c_1, \dots, c_T$ and consider a family of dynamic risk measures $\{\rho_{t,T}^{\pi}\}_{t=1,\dots,T}^{\pi\in\Pi}$ which is translation-invariant, stochastically conditionally time-consistent, and Markov. Our objective is to analyze the risk minimization problem:

$$\min_{\pi\in\Pi} v_1^{\pi}(x_1), \quad x_1 \in \mathcal{X}.$$

For this purpose, we introduce the family of *value functions*:

$$v_t^*(h_t) = \inf_{\pi\in\Pi_{t,T}(h_t)} v_t^{\pi}(h_t), \quad t = 1, \dots, T, \quad h_t \in \mathcal{X}^t, \tag{17}$$

where $\Pi_{t,T}(h_t)$ is the set of feasible deterministic policies $\pi = \{\pi_t, \dots, \pi_T\}$. As stated in Theorem 4.4, transition risk mappings $\{\sigma_t\}_{t=1,\dots,T-1}$ exist, such that

$$\begin{aligned} v_t^{\pi}(h_t) = c_t(x_t, \pi_t(h_t)) + \sigma_t(x_t, Q_t(x_t, \pi_t(h_t)), v_{t+1}^{\pi}(h_t, \cdot)), \\ t = 1, \dots, T-1, \ \pi \in \Pi, \ h_t \in \mathcal{X}^t. \end{aligned} \tag{18}$$

Our intention is to prove that the value functions $v_t^*(\cdot)$ are *memoryless*, that is, for all $h_t = (x_1, \dots, x_t)$ and $h_t' = (x_1', \dots, x_t')$ such that $x_t = x_t'$, we have $v_t^*(h_t) = v_t^*(h_t')$. In this case, with a slight abuse of notation, we shall simply write $v_t^*(x_t)$.

In order to formulate the main result of this subsection, we equip the space $\mathcal{P}(\mathcal{X})$ of probability measures on $\mathcal{X}$ with the topology of weak convergence.

Theorem 4.5. *In addition to the general assumptions, we assume the following conditions:*
(i) *The transition kernels $Q_t(\cdot,\cdot)$, $t = 1,\dots,T$, are continuous;*
(ii) *For every lower semicontinuous $v \in \mathcal{V}$ the transition risk mappings $\sigma_t(\cdot,\cdot,v)$, $t = 1,\dots,T$, are lower semi-continuous;*
(iii) *The functions $c_t(\cdot,\cdot)$, $t = 1,\dots,T$, are lower semicontinuous;*
(iv) *The multifunctions $\mathcal{U}_t(\cdot)$, $t = 1,\dots,T$, are measurable, compact-valued, and upper semi-continuous.*
Then the functions v_t^*, $t = 1,\dots,T$, are measurable, memoryless, lower semicontinuous, and satisfy the following dynamic programming equations:

$$v_T^*(x) = \min_{u\in\mathcal{U}_T(x)} c_T(x,u), \quad x \in \mathcal{X},$$
$$v_t^*(x) = \min_{u\in\mathcal{U}_t(x)} \left\{c_t(x,u) + \sigma_t(x, Q_t(x,u), v_{t+1}^*)\right\}, \quad x \in \mathcal{X}, \quad t = T-1,\dots,1.$$

Moreover, an optimal Markov policy $\hat\pi$ exists and satisfies the equations:

$$\hat\pi_T(x) \in \operatorname*{argmin}_{u\in\mathcal{U}_T(x)} c_T(x,u), \quad x \in \mathcal{X},$$
$$\hat\pi_t(x) \in \operatorname*{argmin}_{u\in\mathcal{U}_t(x)} \left\{c_t(x,u) + \sigma_t(x, Q_t(x,u), v_{t+1}^*)\right\}, \quad x \in \mathcal{X}, \quad t = T-1,\dots,1.$$

Proof. We prove the memoryless property of $v_t^*(\cdot)$ and construct the optimal Markov policy by induction backwards in time. For all $h_T \in \mathcal{X}^T$ we have

$$v_T^*(h_T) = \inf_{\pi\in\Pi} c_T(x_T, \pi_T(h_T)) = \inf_{u\in\mathcal{U}_T(x_T)} c_T(x_T,u). \tag{19}$$

Since $c_T(\cdot,\cdot)$ is lower semicontinuous, it is a *normal integrand*, that is, its epigraphical mapping $x \mapsto \{(u,\alpha) \in \mathcal{U}\times\mathbb{R} : c_T(x,u) \le \alpha\}$ is closed-valued and measurable [37, Ex. 14.31]. Due to assumption (iv), the mapping

$$\bar c_T(x,u) = \begin{cases} c_T(x,u) & \text{if } u \in \mathcal{U}_T(x), \\ +\infty & \text{otherwise,} \end{cases}$$

is a normal integrand as well. By virtue of [37, Thm. 14.37], the infimum in (19) is attained and is a measurable function of x_T. Hence, $v_T^*(\cdot)$ is measurable and memoryless. By assumptions (iii) and (iv) and Berge theorem, it is also lower semicontinuous (see, e.g. [3, Thm. 1.4.16]). Moreover, the optimal solution mapping $\Psi_T(x) = \{u \in \mathcal{U}_T(x) : c_T(x,u) = v_T^*(x)\}$ is measurable and has nonempty and closed values. Therefore, a measurable selector $\hat\pi_T$ of Ψ_T exists [27], [3, Thm. 8.1.3].

Suppose $v_{t+1}^*(\cdot)$ is measurable, memoryless, and lower semicontinuous, and Markov decision rules $\{\hat\pi_{t+1},\dots,\hat\pi_T\}$ exist such that

$$v_{t+1}^*(x_{t+1}) = v_{t+1}^{\{\hat\pi_{t+1},\dots,\hat\pi_T\}}(x_{t+1}), \quad \forall\, h_{t+1} \in \mathcal{X}^{t+1}.$$

Then for any $h_t \in \mathcal{X}^t$ we have

$$v_t^*(h_t) = \inf_{\pi\in\Pi} v_t^\pi(h_t) = \inf_{\pi\in\Pi} \left\{c_t(x_t,\pi_t(h_t)) + \sigma_t(x_t, Q_t(x_t,\pi_t(h_t)), v_{t+1}^\pi(h_t,\cdot))\right\}.$$

On the one hand, since $v_{t+1}^\pi(h_t,\cdot) \ge v_{t+1}^*(\cdot)$ and σ_t is non-decreasing with respect to the last argument, we obtain

$$\begin{aligned} v_t^*(h_t) &\ge \inf_{\pi\in\Pi} \left\{c_t(x_t,\pi_t(h_t)) + \sigma_t(x_t, Q_t(x_t,\pi_t(h_t)), v_{t+1}^*)\right\} \\ &= \inf_{u\in\mathcal{U}_t(x_t)} \left\{c_t(x_t,u) + \sigma_t(x_t, Q_t(x_t,u), v_{t+1}^*)\right\}. \end{aligned} \tag{20}$$

By assumptions (i)–(iii), the mapping $(x,u) \mapsto c_t(x,u) + \sigma_t(x, Q_t(x,u), v^*_{t+1})$ is lower semicontinuous. Invoking [37, Thm. 14.37] and assumption (iv) again, exactly as in the case of $t = T$, we conclude that the optimal solution mapping

$$\begin{aligned}\Psi_t(x) &= \Big\{u \in \mathcal{U}_t(x) : c_t(x,u) + \sigma_t(x, Q_t(x,u), v^*_{t+1}) \\ &= \inf_{u \in \mathcal{U}_t(x)} \Big\{c_t(x,u) + \sigma_t(x, Q_t(x,u), v^*_{t+1})\Big\}\Big\}\end{aligned}$$

is measurable and has nonempty and closed values; hence, a measurable selector $\hat{\pi}_t$ of Ψ_t exists [3, Thm. 8.1.3]. Substituting this selector into (20), we obtain

$$v^*_t(h_t) \geq c_t(x_t, \hat{\pi}_t(x_t)) + \sigma_t(x_t, Q_t(x_t, \hat{\pi}_t(x_t)), v^{\{\hat{\pi}_{t+1},\dots,\hat{\pi}_T\}}_{t+1}) = v^{\{\hat{\pi}_t,\dots,\hat{\pi}_T\}}_t(x_t).$$

In the last equation, we used (18) and the fact that the decision rules $\hat{\pi}_t, \dots, \hat{\pi}_T$ are Markov.

On the other hand,

$$v^*_t(h_t) = \inf_{\pi \in \Pi} v^\pi_t(h_t) \leq v^{\{\hat{\pi}_t,\dots,\hat{\pi}_T\}}_t(x_t).$$

Therefore, $v^*_t(h_t) = v^{\{\hat{\pi}_t,\dots,\hat{\pi}_T\}}_t(x_t)$ is measurable, memoryless, and

$$\begin{aligned}v^*_t(x_t) &= \min_{u \in \mathcal{U}_t(x_t)} \Big\{c_t(x_t,u) + \sigma_t(x_t, Q_t(x,u), v^*_{t+1})\Big\} \\ &= c_t(x_t, \hat{\pi}_t(x_t)) + \sigma_t(x_t, Q_t(x_t, \hat{\pi}_t(x_t)), v^*_{t+1}).\end{aligned}$$

By assumptions (ii), (iii), (iv), and Berge theorem, $v^*_t(\cdot)$ is lower semicontinuous (see, e.g. [3, Thm. 1.4.16]). This completes the induction step. □

Let us verify the weak lower semicontinuity assumption (ii) of the mean–semideviation transition risk mapping of Example 2.13. To make the mapping Markovian, we assume that the parameter $\varkappa$ depends on x only, that is,

$$\sigma(x,q,v) = \int_{\mathcal{X}} v(s)\, q(ds) + \varkappa(x) \left(\int_{\mathcal{X}} \Big[\Big(v(s) - \int_{\mathcal{X}} v(s')\, q(ds')\Big)_+\Big]^p\, q(ds)\right)^{1/p}. \tag{21}$$

As before, $p \in [1, \infty)$. For simplicity, we skip the subscript t of σ and $\varkappa$.

Lemma 4.6. *Suppose $\varkappa(\cdot)$ is continuous. Then for every lower semicontinuous function v, the mapping $(x,q) \mapsto \sigma(x,q,v)$ in* (21) *is lower semicontinuous.*

Proof. Let $q_k \to q$ weakly and $x_k \to x$. For all $s \in \mathcal{X}$ we have the inequality

$$\begin{aligned}0 &\leq \Big[v(s) - \int_{\mathcal{X}} v(s')\, q(ds')\Big]_+ \\ &\leq \Big[v(s) - \int_{\mathcal{X}} v(s')\, q_k(ds')\Big]_+ + \Big[\int_{\mathcal{X}} v(s')\, q_k(ds') - \int_{\mathcal{X}} v(s')\, q(ds')\Big]_+.\end{aligned}$$

By the triangle inequality for the norm in $\mathcal{L}_p(\mathcal{X}, \mathcal{B}(\mathcal{X}), q_k)$,

$$\begin{aligned}&\left(\int_{\mathcal{X}} \Big[v(s) - \int_{\mathcal{X}} v(s')\, q(ds')\Big]_+^p\, q_k(ds)\right)^{1/p} \\ &\qquad \leq \left(\int_{\mathcal{X}} \Big[v(s) - \int_{\mathcal{X}} v(s')\, q_k(ds')\Big]_+^p\, q_k(ds)\right)^{1/p} + \Big[\int_{\mathcal{X}} v(s')\, q_k(ds') - \int_{\mathcal{X}} v(s')\, q(ds')\Big]_+.\end{aligned}$$

Adding $\int_{\mathcal{X}} v(s)\, q(ds)$ to both sides, we obtain

$$\int_{\mathcal{X}} v(s)\; q(ds) + \left(\int_{\mathcal{X}} \left[v(s) - \int_{\mathcal{X}} v(s')\; q(ds') \right]_+^p q_k(ds) \right)^{1/p}$$
$$\leq \left(\int_{\mathcal{X}} \left[v(s) - \int_{\mathcal{X}} v(s')\; q_k(ds') \right]_+^p q_k(ds) \right)^{1/p} + \max \left[\int_{\mathcal{X}} v(s)\; q_k(ds), \int_{\mathcal{X}} v(s)\; q(ds) \right].$$

By the lower semicontinuity of v and weak convergence of q_k to q, we have

$$\int_{\mathcal{X}} v(s)\; q(ds) \leq \liminf_{k\to\infty} \int_{\mathcal{X}} v(s)\; q_k(ds),$$

that is, for every $\varepsilon > 0$, we can find k_ε such that for all $k \geq k_\varepsilon$

$$\int_{\mathcal{X}} v(s)\; q_k(ds) \geq \int_{\mathcal{X}} v(s)\; q(ds) - \varepsilon.$$

Therefore, for these k we obtain

$$\int_{\mathcal{X}} v(s)\; q(ds) + \left(\int_{\mathcal{X}} \left[v(s) - \int_{\mathcal{X}} v(s')\; q(ds') \right]_+^p q_k(ds) \right)^{1/p}$$
$$\leq \int_{\mathcal{X}} v(s)\; q_k(ds) + \left(\int_{\mathcal{X}} \left[v(s) - \int_{\mathcal{X}} v(s')\; q_k(ds') \right]_+^p q_k(ds) \right)^{1/p} + \varepsilon.$$

Taking the "lim inf" of both sides, and using the weak convergence of q_k to q and the lower semicontinuity of the functions integrated, we conclude that

$$\int_{\mathcal{X}} v(s)\; q(ds) + \left(\int_{\mathcal{X}} \left[v(s) - \int_{\mathcal{X}} v(s')\; q(ds') \right]_+^p q(ds) \right)^{1/p}$$
$$\leq \liminf_{k\to\infty} \left\{ \int_{\mathcal{X}} v(s)\; q_k(ds) + \left(\int_{\mathcal{X}} \left[v(s) - \int_{\mathcal{X}} v(s')\; q_k(ds') \right]_+^p q_k(ds) \right)^{1/p} \right\} + \varepsilon.$$

As $\varepsilon > 0$ was arbitrary, the last relation proves the lower semicontinuity of σ in the case when $\varkappa(x) \equiv 1$. The case of a continuous $\varkappa(x) \in [0, 1]$ can be now easily analyzed by noticing that $\sigma(x, q, v)$ is a convex combination of the expected value, and the risk measure of the last displayed relation:

$$\sigma(x, q, v) = (1 - \varkappa(x)) \int_{\mathcal{X}} v(s)\; q(ds)$$
$$+ \varkappa(x) \left\{ \int_{\mathcal{X}} v(s)\; q(ds) + \left(\int_{\mathcal{X}} \left[v(s) - \int_{\mathcal{X}} v(s')\; q(ds') \right]_+^p q(ds) \right)^{1/p} \right\}.$$

As both components are lower semicontinuous in (x, q), so is their sum. □

5 Application to Partially-Observable Markov Decision Processes

In this section, we apply our previous results to a partially-observable Markov decision process $\{X_t, Y_t\}_{t=1,\ldots,T}$, in which $\{X_t\}_{t=1,\ldots,T}$ is observable and $\{Y_t\}_{t=1,\ldots,T}$ is not. We use the term

"partially-observable Markov decision process" (POMDP) in a more general way than the extant literature, because we consider dynamic risk measures of the cost sequence as the objective of control, rather than just the expected value of the total cost.

5.1 The Model

In order to develop our subsequent theory, it is essential to define the model in a clear and rigorous way. This subsection follows Chapter 5 of [4], which the readers are encouraged to consult for more details.

The state space of the model is defined as $\mathcal{X} \times \mathcal{Y}$ where $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ and $(\mathcal{Y}, \mathcal{B}(\mathcal{Y}))$ are two Polish spaces. From the modeling perspective, $x \in \mathcal{X}$ is the part of the state that we can observe at each step, while $y \in \mathcal{Y}$ is unobservable. The measurable space that we will work with is then given by $(\mathcal{X} \times \mathcal{Y})^T$ endowed with the canonical product σ-field, and we use x_t and y_t to denote the canonical projections at time t. The control space is still given by $\mathcal{U}$, and since only $\mathcal{X}$ is observable, the set of admissible controls at step t is still given by a multifunction $\mathcal{U}_t : \mathcal{X} \rightrightarrows \mathcal{U}$. The transition kernel at time t is now $K_t : \text{graph}(\mathcal{U}_t) \times \mathcal{Y} \to \mathcal{P}(\mathcal{X} \times \mathcal{Y})$; in other words, given that at time t the state is (x, y) and we apply control u, the distribution of the next state is $K_t(\cdot\,|x, y, u)$. Finally, the cost incurred at each stage t is given by a measurable and bounded functional $c_t : \text{graph}(\mathcal{U}_t) \to \mathbb{R}$.

At time t, all the information available for making a decision is $g_t = (x_1, u_1, \cdots, x_t)$ because y_t is unobservable. In the following, the notion of a history will always correspond to this observable history. We use $\mathcal{G}_t$ to denote the space of possible histories g_t.

For this controlled process, a (deterministic) *history-dependent admissible policy* $\pi = (\pi_1, \ldots, \pi_T)$ is a sequence of decision rules $\pi_t : \mathcal{G}_t \to \mathcal{U}$ such that $\pi_t(g_t) \in \mathcal{U}_t(x_t)$ for all $g_t \in \mathcal{G}_t$. We can easily prove by induction on t that for such an admissible policy π, each π_t reduces to a measurable function on $\mathcal{X}^t$, as $u_s = \pi_s(x_1, \ldots, x_s)$ for all $s = 1, \ldots, t-1$. We are still using π_s to denote the decision rule, although it is a different function, formally; it will not lead to any misunderstanding. The set of admissible policies is

$$\Pi = \{\pi = (\pi_1, \ldots, \pi_T) \,|\, \pi_t(x_1, \ldots, x_t) \in \mathcal{U}_t(x_t),\ t = 1, \ldots, T\}.$$

Any policy $\pi \in \Pi$ defines a process $\{X_t, Y_t, U_t\}_{t=1,\ldots,T}$, with $U_t = \pi_t(X_1, \ldots, X_t)$ and the resulting cost process $Z_t = c_t(X_t, U_t)$, $t = 1, \ldots, T$. Our objective is to define stochastically time consistent risk measures for this problem and to optimize them with respect to the policy $\pi \in \Pi$. To this end, we shall convert the problem to a problem of controlling the observable part of the process, $\{X_t\}_{t=1,\ldots,T}$, with history-dependent transition kernels. This can be accomplished with the use of the Bayes operator.

5.2 Bayes Operator

In a POMDP, the Bayes operator provides a way to update from prior belief to posterior belief. More precisely, assume that our current state observation is x and action is u, and our current best guess of the distribution of the unobservable state is ξ. Given a new observation x', we can find a formula to determine the posterior distribution of the unobservable state.

Let us start with a fairly general construction of the Bayes operator. Assuming the above setup, for given $(x, \xi, u) \in \mathcal{X} \times \mathcal{P}(\mathcal{Y}) \times \mathcal{U}$, define a new measure $m_t(x, \xi, u)$ on $\mathcal{X} \times \mathcal{Y}$, initially on all measurable rectangles $A \times B$ as

$$m_t(x, \xi, u)(A \times B) = \int_{\mathcal{Y}} K_t(A \times B \,|\, x, y, u)\, \xi(dy).$$

We verify readily that this uniquely defines a probability measure on $\mathcal{X} \times \mathcal{Y}$. If the measurable space $(\mathcal{Y}, \mathcal{B}(\mathcal{Y}))$ is standard, *i.e.*, isomorphic to a Borel subspace of $\mathbb{R}$, we can disintegrate $m_t(x, \xi, u)$ into its marginal $\lambda_t(x, \xi, u)(dx')$ on $\mathcal{X}$ and a transition kernel $\Upsilon_t(x, \xi, u)(x', dy')$ from $\mathcal{X}$ to $\mathcal{Y}$:

$$m_t(x, \xi, u)(dx', dy') = \lambda_t(x, \xi, u)(dx')\, \Upsilon_t(x, \xi, u)(x', dy').$$

For all $C \in \mathcal{B}(\mathcal{Y})$, we define the *Bayes operator* of the POMDP as follows:

$$\Phi_t(x, \xi, u, x')(C) = \Upsilon_t(x, \xi, u)(x', C).$$

The above argument shows that the the Bayes operator exists and is unique as long as the space $\mathcal{Y}$ is standard, which is almost always the case in applications of POMDP. In the following consierations, we always assume that the Bayes operator exists.

Example 5.1 (Bayes operator with kernels given by density functions)**.** *Assume that each transition kernel $K_t(x, y, u)$ has a density $q_t(\cdot, \cdot \mid x, y, u)$ with respect to a finite product measure $\mu_X \otimes \mu_Y$ on $\mathcal{X} \times \mathcal{Y}$. Then the Bayes operator has the form*

$$[\Phi_t(x, \xi, u, x')](A) = \frac{\int_A \int_{\mathcal{Y}} q_t(x', y' \mid x, y, u)\, \xi(dy)\, \mu_Y(dy')}{\int_{\mathcal{Y}} \int_{\mathcal{Y}} q_t(x', y' \mid x, y, u)\, \xi(dy)\, \mu_Y(dy')}, \quad \forall\, A \in \mathcal{B}(\mathcal{Y}). \tag{22}$$

In particular, if $\mathcal{Y}$ is a finite space, then

$$[\Phi_t(x, \xi, u, x')](\{y'\}) = \frac{\sum_{y \in \mathcal{Y}} q_t(x', y' \mid x, y, u)\, \xi(\{y\})}{\sum_{z \in \mathcal{Y}} \sum_{y \in \mathcal{Y}} q_t(x', z \mid x, y, u)\, \xi(\{y\})}, \quad y' \in \mathcal{Y}.$$

If the formulas above have a zero denominator for some (x, ξ, u, x'), we can formally define $\Phi_t(x, \xi, u, x')$ to be an arbitrarily selected distribution on $\mathcal{Y}$.

5.3 Risk Measures Based on the Observable Part of the Process

If $\Xi_1 : \mathcal{X} \to \mathcal{P}(\mathcal{Y})$ is the conditional distribution of Y_1, given X_1, we can consider POMDP as a special case of a fully-observable non-Markov model discussed in Section 3.1, with the process X (and omitting Y), because the admissible control sets $\mathcal{U}_t(x_t)$, the costs $c_t(x_t, u_t)$, and the kernel $Q_t(dX_{t+1} | x_1, u_1, \ldots, x_t, u_t)$ do not depend on Y. To justify this transformation, we make the following observations, valid for all $t = 1, \ldots, T$:

1. The control $u_t \in \mathcal{U}_t(x_t)$ can be applied after observing x_t.

2. After observing x_t, with the help of the Bayes operator, the conditional distribution of Y_t is a function of $g_t = (x_1, u_1, \ldots, x_{t-1}, u_{t-1}, x_t)$ and can be calculated iteratively with $\Xi_1(g_1) = \Xi_1(x_1)$ and

$$\Xi_t(g_t) = \Phi_{t-1}\Big(x_{t-1}, \Xi_{t-1}(g_{t-1}), u_{t-1}, x_t\Big), \quad t = 2, \ldots, T. \tag{23}$$

3. For $\xi_t = \Xi_t(g_t)$, after applying a control u_t, the conditional distribution of X_{t+1} is given by the integral:

$$Q_t(x_t, \xi_t, u_t) = \int_{y \in \mathcal{Y}} K_t^X(x_t, y, u_t)\, \xi_t(dy), \tag{24}$$

where $K_t^X(x_t, y, u_t)$ is the marginal distribution of $K_t(x_t, y, u_t)$ on $\mathcal{X}$. The precise meaning (*e.g.*, set-wise or weak) of the integral depends on the selected topology in the space $\mathcal{P}(\mathcal{X})$.

4. We can define an admissible decision rule as a measurable selector $\pi_t : g_t \mapsto \mathcal{U}_t(x_t)$, because g_t encompasses all information available before making a decision at time t. As we discussed it in Section 3.1, for an admissible policy $\pi = (\pi_1, \dots, \pi_T)$, each π_t reduces to a function of $h_t = (x_1, \dots, x_t)$, and the set of admissible policies is as in (11), but with a reduced form of $\mathcal{U}_t$, as in (22). For each $\pi \in \Pi$, we denote

$$\Xi_t^\pi(h_t) = \Xi_t(x_1, \pi_1(x_1), \dots, x_{t-1}, \pi_{t-1}(x_1, \dots, x_{t-1}), x_t)$$

and

$$Q_t^\pi(h_t) = Q_t(x_t, \Xi_t^\pi(h_t), \pi_t(h_t)).$$

5. For any $\pi \in \Pi$, the function $h_t \mapsto c_t(x_t, \pi_t(h_t))$ is an element of $\mathcal{Z}_t$, the exact same space that we defined in (2) and used in section 3.

The integral in (24) is a subtle issue, especially for the dynamic programming part, where semicontinuity properties of various elements of the model are relevant. To avoid serious technical complications, we assume from now on that the unobserved state space $\mathcal{Y}$ is finite, and thus all convergence concepts in (24) become equivalent.

Therefore, without any model limitation or information loss at each time t, we can simply consider a POMDP as a system with the observable process X only, whose dynamics is described in (24). Consequently, we can construct dynamic risk measures as in Sections 2 and 3. To alleviate notation, for such family of process-based dynamic risk measures $\{\rho_{t,T}^\pi\}_{t=1,\dots,T}^{\pi\in\Pi}$, for all $\pi \in \Pi$, and for all measurable $c = (c_1, \dots, c_T)$, we still write

$$v_t^{c,\pi}(h_t) := \rho_{t,T}^\pi(c_t(X_t, \pi_t(H_t)), \dots, c_T(X_T, \pi_T(H_T)))(h_t). \tag{25}$$

The following statement specializes Corollary 4.1 to the partially observable case.

Corollary 5.2. *For a POMDP, a family of process-based dynamic risk measures $\{\rho_{t,T}^\pi\}_{t=1,\dots,T}^{\pi\in\Pi}$ is translation-invariant and stochastically conditionally time-consistent if and only if there exist functionals*

$$\sigma_t : \Big\{(h_t, Q_t^\pi(h_t)) : \pi \in \Pi,\ h_t \in \mathcal{X}^t\Big\} \times \mathcal{V} \to \mathbb{R}, \quad t = 1 \dots T-1,$$

in which $\mathcal{V}$ is the set of measurable and bounded functions on $\mathcal{X}$, such that:
(i) *For all $t = 1, \dots, T-1$, and all $h_t \in \mathcal{X}^t$, the mapping $\sigma(h_t, \cdot, \cdot)$ is normalized and strongly monotonic with respect to stochastic dominance on $\{\, Q_t^\pi(h_t) : \pi \in \Pi \,\}$;*
(ii) *For all $\pi \in \Pi$, for all bounded measurable c, for all $t = 1, \dots, T-1$, and for all $h_t \in \mathcal{X}^t$,*

$$v_t^{c,\pi}(h_t) = c_t(x_t, \pi_t(h_t)) + \sigma_t(h_t, Q_t^\pi(h_t), v_{t+1}^{c,\pi}(h_t, \cdot)). \tag{26}$$

5.4 Markov Property

For finite $\mathcal{Y}$, the set $\mathcal{P}(\mathcal{Y})$ is the simplex $\Big\{\xi \in \mathbb{R}^{|\mathcal{Y}|} : \xi \geq 0,\ |\xi|_1 = 1\Big\}$; obviously, a Borel subset of a finite dimensional real space. We can deduce from (23) and (24) that $(X_t, \Xi_t)_{t=1,\dots,T}$ is a Markov process with the following transition rules:

$$\begin{cases} [Q_t(x_t, \xi_t, u_t)](dx') = \displaystyle\sum_{y\in\mathcal{Y}} \Big[K_t^X(x_t, y, u_t)\Big](dx')\,\xi_t(\{y\}), \\ \Xi_{t+1}(g_{t+1}) = \Phi_t\Big(x_t, \Xi_t(g_t), u_t, x_{t+1}\Big). \end{cases} \tag{27}$$

We say that an admissible policy is *Markov* if each decision rule depends only on the current observed state x and the current belief state ξ.

Definition 5.3. *In POMDP, a policy $\pi \in \Pi$ is* Markov *if $\pi_t(h_t) = \pi_t(h'_t)$ for all $t = 1, \dots, T$ and all h_t, $h'_t \in \mathcal{X}^t$ such that $x_t = x'_t$ and $\Xi_t^\pi(h_t) = \Xi_t^\pi(h'_t)$.*

For a fixed Markov policy π, the future evolution of the process $\{(X_\tau, \Xi_\tau^\pi)\}_{\tau=t,\dots,T}$ is solely dependent on the current $(x_t, \Xi_t^\pi(h_t))$, and so is the distribution of the future costs $c_\tau(X_\tau, \pi_\tau(X_\tau))$, $\tau = t, \dots, T$. Therefore, we can define the Markov property of risk measures for POMDP.

Definition 5.4. *A family of process-based dynamic risk measures $\{\rho_{t,T}^\pi\}_{t=1,\dots,T}^{\pi\in\Pi}$ for a POMDP is* Markov *if for all Markov policies $\pi \in \Pi$, for all measurable $c = (c_1, \dots, c_T)$, and for all $h_t = (x_1, \dots, x_t)$ and $h'_t = (x'_1, \dots, x'_t)$ in $\mathcal{X}^t$ such that $x_t = x'_t$ and $\Xi_t^\pi(h_t) = \Xi_t^\pi(h'_t)$, we have*

$$v_t^{c,\pi}(h_t) = v_t^{c,\pi}(h'_t).$$

Proposition 5.5. *A translation invariant and stochastically conditionally time-consistent family of process-based dynamic risk measures $\{\rho_{t,T}^\pi\}_{t=1,\dots,T}^{\pi\in\Pi}$ is Markov if and only if the dependence of σ_t on h_t is carried by $(x_t, \Xi_t^\pi(h_t))$ only, for all $t = 1, \dots, T-1$.*

Proof. A minor modification of the proof of Proposition 4.3 is sufficient. □

The following theorem summarizes our observations.

Theorem 5.6. *A family of process-based dynamic risk measures $\{\rho_{t,T}^\pi\}_{t=1,\dots,T}^{\pi\in\Pi}$ for a POMDP is translation-invariant, stochastically conditionally time-consistent and Markov if and only if there exist functionals*

$$\sigma_t : \left\{(x_t, \Xi_t^\pi(h_t), Q_t^\pi(h_t)) : \pi \in \Pi,\ h_t \in \mathcal{X}^t\right\} \times \mathcal{V} \to \mathbb{R}, \quad t = 1 \dots T-1,$$

such that
(i) *for all $t = 1, \dots, T-1$ and all $(x, \xi) \in \{(x_t, \Xi_t^\pi(h_t)) : \pi \in \Pi,\ h_t \in \mathcal{X}^t\}$, $\sigma_t(x, \xi, \cdot, \cdot)$ is normalized and strongly monotonic with respect to stochastic dominance on $\{\, Q_t^\pi(h_t) : \pi \in \Pi,\ h_t \in \mathcal{X}^t$ such that $x_t = x,\ \Xi_t^\pi(h_t) = \xi \,\}$;*
(ii) *for all $\pi \in \Pi$, for all measurable bounded c, for all $t = 1, \dots, T-1$, and for all $h_t \in \mathcal{X}^t$,*

$$v_t^{c,\pi}(h_t) = c_t(x_t, \pi_t(h_t)) + \sigma_t(x_t, \Xi_t^\pi(h_t), Q_t^\pi(h_t), v_{t+1}^{c,\pi}(h_t, \cdot)). \tag{28}$$

This allows us to evaluate risk of Markov policies in a recursive way.

Corollary 5.7. *Under the conditions of Theorem 5.6, for any Markov policy π, the function* (25) *depends on c, $\pi_t, \dots, \pi_T$, and (x_t, ξ_t) only, and the following relation is true:*

$$v_t^{c,\pi_t,\dots,\pi_T}(x_t, \xi_t) = c_t(x_t, \pi_t(x_t, \xi_t)) + \sigma_t(x_t, \xi_t, Q_t(x_t, \xi_t, \pi_t(x_t, \xi_t)), x' \mapsto v_{t+1}^{c,\pi_{t+1},\dots,\pi_T}(x', \Phi_t(x_t, \xi_t, \pi_t(x_t, \xi_t), x'))). \tag{29}$$

Proof. We prove the relation by induction backward in time. For $t = T$ we have $v_T^{c,\pi}(h_T) = c_T(x_T, \pi_T(x_T, \xi_T))$ and our assertion is true. If it is true for $t+1$, formula (28) reads

$$v_t^{c,\pi}(h_t) = c_t(x_t, \pi_t(x_t, \xi_t)) + \sigma_t(x_t, \xi_t, Q_t(x_t, \xi_t, \pi_t(x_t, \xi_t)), x' \mapsto v_{t+1}^{c,\pi_{t+1},\dots,\pi_T}(x', \Phi_t(x_t, \xi_t, \pi_t(x_t, \xi_t), x'))),$$

which proves our assertion. □

5.5 Dynamic Programming

As in subsection 4.2, we fix the cost functions $c_1, \dots, c_T$ and consider a family of dynamic risk measures $\{\rho^{\pi}_{t,T}\}^{\pi\in\Pi}_{t=1,\dots,T}$ which is translation-invariant, stochastically conditionally time-consistent and Markov. Our objective is to analyze the risk minimization problem:

$$\min_{\pi\in\Pi} v^{\pi}_1(x_1, \xi_1), \quad x_1 \in \mathcal{X},\ \xi_1 \in \mathcal{P}(\mathcal{X}).$$

For this purpose, we introduce the family of *value functions*:

$$v^*_t(g_t) = \inf_{\pi\in\Pi_{t,T}(g_t)} v^{\pi}_t(h_t), \quad t = 1, \dots, T, \quad g_t \in \mathcal{G}_t, \tag{30}$$

where $\Pi_{t,T}(g_t)$ is the set of feasible deterministic policies $\pi = \{\pi_1, \dots, \pi_T\}$ such that $\pi_\tau(h_\tau) = u_\tau$ for all $\tau = 1, \dots, t-1$. At this stage of our analysis, we must allow the value functions to depend on the entire history g_t involving past controls, because it is necessary for the reconstruction of the kernels $Q^{\pi}_t, \dots, Q^{\pi}_{T-1}$. As stated in Theorem 5.6, there exist transition risk mappings $\{\sigma_t\}_{t=1,\dots,T-1}$ such that equations (28) hold.

We assume that the space $\mathcal{P}(\mathcal{X})$ is equipped with the topology of weak convergence, and the space $\mathcal{V}$ is equipped with the topology of pointwise convergence. Recall that the space $\mathcal{Y}$ is assumed finite, and thus all convergence concepts in $\mathcal{P}(\mathcal{Y})$ coincide. All continuity statements are made with respect to the said topologies.

We also assume that the kernels $K_t(x, y, u)$ have densities $q_t(\cdot, \cdot \mid x, y, u)$ with respect to a finite product measure $\mu_X \otimes \mu_Y$ on $\mathcal{X} \times \mathcal{Y}$, as in Example 5.1. In this case,

$$\left[\int_{\mathcal{Y}} K^X_t(x, \cdot, u)\, d\xi\right](dx') = \left[\sum_{y'\in\mathcal{Y}} \sum_{y\in\mathcal{Y}} q_t(x', y' \mid x, y, u)\, \xi(\{y\})\right] \mu_X(dx'). \tag{31}$$

The following theorem is an analogue of Theorem 4.5. Its main result is that the value functions (30) are *memoryless*, that is, they depend on (x_t, ξ_t) only, and that they satisfy a generalized form of a dynamic programming equation. The equation also allows us to identify the optimal policy.

Theorem 5.8. *In addition to the general assumptions, we assume the following conditions:*
(i) *The functions $(x, u) \mapsto q_t(x', y'|x, y, u)$ are uniformly bounded and continuous;*
(ii) *The transition risk mappings $\sigma_t(\cdot, \cdot, \cdot, \cdot)$, $t = 1, \dots, T$, are lower semi-continuous;*
(iii) *The functions $c_t(\cdot, \cdot)$, $t = 1, \dots, T$, are lower semicontinuous;*
(iv) *The multifunctions $\mathcal{U}_t(\cdot)$, $t = 1, \dots, T$, are compact-valued and upper-semicontinuous.*
Then the functions v^*_t, $t = 1, \dots, T$ are memoryless, lower semicontinuous, and satisfy the following dynamic programming equations:

$$v^*_T(x, \xi) = \min_{u\in\mathcal{U}_T(x)} c_T(x, u), \quad x \in \mathcal{X}, \quad \xi \in \mathcal{P}(\mathcal{X}), \tag{32}$$

$$\begin{aligned} v^*_t(x, \xi) = \min_{u\in\mathcal{U}_t(x)} \Big\{ & c_t(x, u) + \\ & \sigma_t\Big(x, \xi, \int_{\mathcal{Y}} K^X_t(x, \cdot, u)\, d\xi, x' \mapsto v^*_{t+1}(x', \Phi_t(x, \xi, u, x'))\Big)\Big\}, \\ & x \in \mathcal{X}, \quad \xi \in \mathcal{P}(\mathcal{Y}), \quad t = T-1, \dots, 1. \end{aligned} \tag{33}$$

Moreover, an optimal Markov policy $\hat{\pi}$ exists and satisfies the equations:

$$\hat{\pi}_T(x,\xi) \in \operatorname*{argmin}_{u\in\mathcal{U}_T(x)} c_T(x,u), \quad x\in\mathcal{X}, \quad \xi\in\mathcal{P}(\mathcal{Y}), \tag{34}$$

$$\begin{aligned}\hat{\pi}_t(x,\xi) \in \operatorname*{argmin}_{u\in\mathcal{U}_t(x)} \Big\{ & c_t(x,u) + \\ & \sigma_t\Big(x,\xi,\int_{\mathcal{Y}} K_t^X(x,\cdot,u)\,d\xi, x'\mapsto v_{t+1}^*(x',\Phi_t(x,\xi,u,x'))\Big)\Big\}, \\ & x\in\mathcal{X}, \quad \xi\in\mathcal{P}(\mathcal{Y}), \quad t=T-1,\dots,1.\end{aligned} \tag{35}$$

Proof. We follow the pattern of the proof of Theorem 4.5, with adjustments and refinements due to the specific structure of the partially observable case. Recall that $\mathcal{G}_t$ is the space of observable state and control histories up to time t; we denote its generic elements as $g_t = (x_1,u_1,\dots,u_{t-1},x_t)$. For all $g_T\in\mathcal{G}_T$ we have

$$v_T^*(g_T) = \inf_{\pi\in\Pi_{T,T}(g_T)} c_T(x_T,\pi_T(h_T)) = \inf_{\pi_T} c_T(x_T,\pi_T(h_T)) = \inf_{u\in\mathcal{U}_T(x_T)} c_T(x_T,u). \tag{36}$$

By assumptions (iii) and (iv), owing to the Berge theorem (see [3, Theorem 1.4.16]), the infimum in (36) is attained and is a lower semicontinuous function of x_T. Hence, v_T^* is memoryless. Moreover, the optimal solution mapping $\Psi_T(x) = \{u\in\mathcal{U}_T(x) : c_T(x,u) = v_T^*(x)\}$ has nonempty and closed values and is measurable. Therefore, a measurable selector $\hat{\pi}_T$ of Ψ_T exists (see, [27], [3, Thm. 8.1.3]).

Suppose $v_{t+1}^*(\cdot)$ is memoryless, lower semicontinuous, and Markov decision rules $\{\hat{\pi}_{t+1},\dots,\hat{\pi}_T\}$ exist such that, by virtue of Corollary 5.7,

$$v_{t+1}^*(x_{t+1},\xi_{t+1}) = v_{t+1}^{\{\hat{\pi}_{t+1},\dots,\hat{\pi}_T\}}(x_{t+1},\xi_{t+1}), \quad \forall h_{t+1}\in\mathcal{X}^{t+1}.$$

Then for any $h_t\in\mathcal{X}^t$ formula (28) yields

$$\begin{aligned} v_t^*(g_t) &= \inf_{\pi\in\Pi_{t,T}(g_t)} v_t^\pi(h_t) \\ &= \inf_{\pi\in\Pi_{t,T}(g_t)} \left\{ c_t(x_t,\pi_t(h_t)) + \sigma_t\Big(x_t,\xi_t,\int_{\mathcal{Y}} K_t^X(x_t,\cdot,\pi_t(h_t))\,d\xi_t, v_{t+1}^\pi(h_t,\cdot)\Big)\right\}. \end{aligned}$$

Since $v_{t+1}^\pi(h_t,x') \ge v_{t+1}^*(x',\Phi_t(x_t,\xi_t,\pi_t(h_t),x'))$ for all $x'\in\mathcal{X}$, and σ_t is non-decreasing with respect to the last argument, we obtain

$$\begin{aligned} v_t^*(g_t) &\ge \inf_{\pi\in\Pi_{t,T}(g_t)} \Big\{ c_t(x_t,\pi_t(h_t)) \\ &\qquad + \sigma_t\Big(x_t,\xi_t,\int_{\mathcal{Y}} K_t^X(x_t,\cdot,\pi_t(h_t))\,d\xi_t, x'\mapsto v_{t+1}^*(x',\Phi_t(x_t,\xi_t,\pi_t(h_t),x'))\Big)\Big\} \\ &= \inf_{u\in\mathcal{U}_t(x_t)} \Big\{ c_t(x_t,u) + \sigma_t\Big(x_t,\xi_t,\int_{\mathcal{Y}} K_t^X(x_t,\cdot,u)\,d\xi_t, x'\mapsto v_{t+1}^*(x',\Phi_t(x_t,\xi_t,u,x'))\Big)\Big\}. \end{aligned} \tag{37}$$

In order to complete the induction step, we need to establish lower semicontinuity of the mapping

$$(x,\xi,u)\mapsto \sigma_t\Big(x,\xi,\int_{\mathcal{Y}} K_t^X(x,\cdot,u)\,d\xi, x'\mapsto v_{t+1}^*(x',\Phi_t(x,\xi,u,x'))\Big). \tag{38}$$

To this end, suppose $x^{(k)}\to x$, $\xi^{(k)}\to\xi$, $u^{(k)}\to u$, as $k\to\infty$.

First, we verify that the mapping $(x,\xi,u)\mapsto\int_{\mathcal{Y}} K_t^X(x,\cdot,u)\,d\xi$ appearing in the third argument of σ_t is continuous. By formula (31), assumption (i), and Lebesgue theorem, for any bounded continuous function $f:\mathcal{X}\to\mathbb{R}$ we have

$$\begin{aligned}\int_{\mathcal{X}} f(x')\left[\int_{\mathcal{Y}} K_t^X(x^{(k)},\cdot,u^{(k)})\,d\xi^{(k)}\right](dx')&\\ =\int_{\mathcal{X}} f(x')\left[\sum_{y'\in\mathcal{Y}}\sum_{y\in\mathcal{Y}} q_t(x',y'\mid x^{(k)},y,u^{(k)})\,\xi^{(k)}(\{y\})\right]\mu_X(dx')&\\ \to\int_{\mathcal{X}} f(x')\left[\sum_{y'\in\mathcal{Y}}\sum_{y\in\mathcal{Y}} q_t(x',y'\mid x,y,u)\,\xi(\{y\})\right]\mu_X(dx')&\\ =\int_{\mathcal{X}} f(x')\left[\int_{\mathcal{Y}} K_t^X(x,\cdot,u)\,d\xi\right](dx').&\end{aligned}$$

Thus, the third argument of σ_t in (38) is continuous with respect to (x,ξ,u).

Let us examine the last argument of σ_t in (38). Owing to assumption (i), the operator

$$[\Phi_t(x,\xi,u,x')](\{y'\})=\frac{\sum_{y\in\mathcal{Y}} q_t(x',y'\mid x,y,u)\,\xi(\{y\})}{\sum_{z\in\mathcal{Y}}\sum_{y\in\mathcal{Y}} q_t(x',z\mid x,y,u)\,\xi(\{y\})},\quad y'\in\mathcal{Y}.$$

is continuous with respect to (x,ξ,u) at all points (x,ξ,u,x') at which

$$\sum_{z\in\mathcal{Y}}\sum_{y\in\mathcal{Y}} q_t(x',z\mid x,y,u)\,\xi(\{y\})>0. \tag{39}$$

Consider the sequence of functions $V^{(k)}:\mathcal{X}\to\mathbb{R}$, $k=1,2,\ldots$, and the function $V:\mathcal{X}\to\mathbb{R}$, defined for $x'\in\mathcal{X}$ as follows:

$$\begin{aligned}V^{(k)}(x')&=v_{t+1}^*(x',\Phi_t(x^{(k)},\xi^{(k)},u^{(k)},x')),\\ V(x')&=v_{t+1}^*(x',\Phi_t(x,\xi,u,x')).\end{aligned}$$

Since $v_{t+1}^*(\cdot,\cdot)$ is lower-semicontinuous and $\Phi_t(\cdot,\cdot,\cdot,x')$ is continuous, whenever condition (39) is satisfied, we infer that

$$V(x')\le\liminf_{k\to\infty} V^{(k)}(x'),$$

for all $x'\in\mathcal{X}$ satisfying (39).

As v_{t+1}^* and Φ_t are measurable, both V and $\liminf_{k\to\infty} V^{(k)}$ are measurable as well. By the fact that the mapping σ_t is preserving the stochastic order $\preceq_{\mathrm{st}}$ of the last argument with respect to the measure $\int_{\mathcal{Y}} K_t^X(x,\cdot,u)\,d\xi$, and by assumption (ii), with the view at the already established continuity of the third argument, we obtain the following chain of relations:

$$\begin{aligned}\sigma_t\Big(x,\xi,\int_{\mathcal{Y}} K_t^X(x,\cdot,u)\,d\xi,V\Big)&\le\sigma_t\Big(x,\xi,\int_{\mathcal{Y}} K_t^X(x,\cdot,u)\,d\xi,\liminf_{k\to\infty} V^{(k)}\Big)\\ &=\sigma_t\Big(x,\xi,\lim_{k\to\infty}\int_{\mathcal{Y}} K_t^X(x^{(k)},\cdot,u^{(k)})\,d\xi^{(k)},\liminf_{k\to\infty} V^{(k)}\Big)\\ &\le\liminf_{k\to\infty}\sigma_t\Big(x^{(k)},\xi^{(k)},\int_{\mathcal{Y}} K_t^X(x^{(k)},\cdot,u^{(k)})\,d\xi^{(k)},V^{(k)}\Big).\end{aligned}$$

Consequently, the mapping (38) is lower semicontinuous.

Using assumptions (ii) and (iv) and invoking Berge theorem again (see, e.g., [3, Theorem 1.4.16]), we deduce that the infimum in (37) is attained and is a lower semicontinuous function of (x_t, ξ_t). Moreover, the optimal solution mapping, that is, the set of $u \in \mathcal{U}_T(x)$ at which the infimum in (37) is attained, is nonempty, closed-valued, and measurable. Therefore, a minimizer $\hat{\pi}_t$ in (37) exists and is a measurable function of (x_t, ξ_t) (see, *e.g.*, [27], [3, Thm. 8.1.3]). Substituting this minimizer into (37), we obtain

$$\begin{aligned} v_t^*(g_t) &\geq c_t(x_t, \hat{\pi}_t(x_t, \xi_t)) \\ &\quad + \sigma_t\Big(x_t, \xi_t, \int_{\mathcal{Y}} K_t^X(x_t, \cdot, \hat{\pi}_t(x_t, \xi_t))\, d\xi_t, x' \mapsto v_{t+1}^*(x', \Phi_t(x, \xi, \hat{\pi}_t(x_t, \xi_t), x'))\Big) \\ &= v_t^{\{\hat{\pi}_t, \ldots, \hat{\pi}_T\}}(x_t, \xi_t). \end{aligned}$$

In the last equation, we used Corollary 5.4. On the other hand, we have

$$v_t^*(g_t) = \inf_{\pi \in \Pi_{t,T}} v_t^{\pi}(h_t) \leq v_t^{\{\hat{\pi}_t, \ldots, \hat{\pi}_T\}}(x_t, \xi_t).$$

Therefore $v_t^*(g_t) = v_t^{\{\hat{\pi}_t, \ldots, \hat{\pi}_T\}}(x_t, \xi_t)$ is memoryless, lower semicontinuous, and

$$\begin{aligned} &v_t^*(x_t, \xi_t) \\ &= \min_{u \in \mathcal{U}_t(x_t)} \left\{ c_t(x_t, u) + \sigma_t\Big(x_t, \xi_t, \int_{\mathcal{Y}} K_t^X(x_t, \cdot, u)\, d\xi_t, x' \mapsto v_{t+1}^*(x', \Phi_t(x_t, \xi_t, u, x'))\Big) \right\} \\ &= c_t(x_t, \hat{\pi}_t(x_t, \xi_t)) \\ &\quad + \sigma_t\Big(x_t, \xi_t, \int_{\mathcal{Y}} K_t^X(x_t, \cdot, \hat{\pi}_t(x_t, \xi_t))\, d\xi_t, x' \mapsto v_{t+1}^*(x', \Phi_t(x_t, \xi_t, \hat{\pi}_t(x_t, \xi_t), x'))\Big). \end{aligned}$$

This completes the induction step. □

We can remark here that Theorem 5.8 remains valid with more general bounded, measurable, and lower semicontinuous cost functions $c_t : \mathcal{X} \times \mathcal{P}(\mathcal{Y}) \times \mathcal{U} \to \mathbb{R}$, which explicitly depend on the belief state ξ_t. The formulation is almost the same, just the cost part in formulas (32)–(35) becomes $c_t(x, \xi, u)$. Formally, we could have also made the sets $\mathcal{U}_t$ dependent on ξ_t, but this is hard to justify from the practical point of view.

6 Illustration: Machine Deterioration

6.1 Description of the process

We consider the problem of minimizing costs of using a machine in T periods. The condition of the machine is deteriorating over time, but is not known with certainty. The only information available is the operating cost. The control in any period is to continue using the machine, or to replace it.

At the beginning of period $t = 1, \ldots, T$, the condition of the machine is denoted by $y_t \in \{1, 2\}$, with 1 denoting the "good" state, and 2 the "bad" state. The controls are denoted by $u_t \in \{0, 1\}$, with 0 meaning "continue", and 1 meaning "replace".

The dynamics is Markovian, with the following transition graph for the "continue" control:

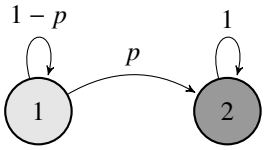

and the following transition matrices $\{K^{[u]}\}$, $u \in \{0, 1\}$:

$$K^{[0]} = \begin{pmatrix} 1-p & p \\ 0 & 1 \end{pmatrix}, \quad K^{[1]} = \begin{pmatrix} 1-p & p \\ 1-p & p \end{pmatrix}. \tag{40}$$

The condition of the machine y_t cannot be observed, but we can observe the cost incurred during the period t, denoted by x_{t+1}. The increment of the time index is due to the fact that this cost becomes known at the end of the period, and provides information for the decision making in the next period. The conditional distribution of x_{t+1}, given y_t, is described by two density functions f_1 and f_2, *i.e.*,

$$\forall C \in \mathbb{R}, \quad \begin{cases} \mathbb{P}[x_{t+1} \le C \,\big|\, y_t = i, u_t = 0] = \displaystyle\int_{-\infty}^{C} f_i, \quad i = 1, 2, \\ \mathbb{P}[x_{t+1} \le C \,\big|\, y_t = i, u_t = 1] = \displaystyle\int_{-\infty}^{C} f_1, \quad i = 1, 2. \end{cases} \tag{41}$$

Assumption 6.1. *The functions f_1 and f_2 are uniformly bounded and the conditional distribution of x_{t+1} given that the machine is in "good" condition is stochastically smaller than the conditional distribution of x_{t+1} given that the machine is in "bad" condition,* i.e.*,*

$$\int_{-\infty}^{C} f_1(x)\, dx \ge \int_{-\infty}^{C} f_2(x)\, dx, \quad \forall C \in \mathbb{R};$$

with a slight abuse of notation, we write it $f_1 \preceq_{\mathrm{st}} f_2$.

Thus the relations (40) and (41) define $(x_t, y_t)_{t=1,\dots,T+1}$ as a partially observable Markov process controlled by $(u_t)_{t=1,\dots,T}$. Based on observations $(x_1, \dots, x_t)$, the belief state $\xi_t \in [0, 1]$ denotes the conditional probability that $y_t = 1$, and we can update the posterior belief state

$$\xi_{t+1} = \begin{cases} \Phi(\xi_t, x_{t+1}), & \text{if } u_t = 0; \\ 1 - p, & \text{if } u_t = 1. \end{cases} \tag{42}$$

where Φ is the Bayes operator,

$$\Phi(\xi, x') = \frac{(1-p)\xi f_1(x')}{\xi f_1(x') + (1-\xi) f_2(x')}. \tag{43}$$

We assume that the probability $\xi_0 \in [0, 1]$ that initially the machine is good, is known; then $\xi_1(x_1) = \Phi(\xi_0, x_1)$. Directly from (43) we see that $\Phi(0, \cdot) = 0$, $\Phi(1, \cdot) = 1 - p$, and $\Phi(\cdot, x')$ is non-decreasing.

6.2 Risk modeling

At the beginning of the period t, if we replace the machine ($u_t = 1$), there is an additional fixed replacement cost R. Then the costs incurred are

$$\begin{cases} c_t(x_t, u_t) = R \cdot u_t + x_t, \quad t = 1, \dots, T; \\ c_{T+1}(x_{T+1}) = x_{T+1}. \end{cases} \tag{44}$$

We denote the history of observations by $h_t = (x_1, \dots, x_t)$ and the set of all history-dependent policies by $\Pi := \{\pi = (\pi_1, \dots, \pi_T) \,|\, \forall t,\ \pi_t(x_1, \dots, x_t) \in \{0, 1\}\}$. We want to evaluate the costs (44) for any $\pi \in \Pi$, and find an optimal policy. The standard risk-neutral approach is to evaluate the conditional expectations of the sum of future costs:

$$\mathbb{E}\left[\sum_{\tau=t}^{T} (R \cdot \pi_\tau(h_\tau) + x_\tau) + x_{t+1} \,\middle|\, h_t\right], \quad h_t \in \mathbb{R}^t, \quad t = 1, \dots, T+1, \quad \pi \in \Pi.$$

As shown in Section 5, we can construct risk measures $\{v_t^\pi : \mathbb{R}^t \to \mathbb{R}\}_{t=1,\dots,T+1}^{\pi \in \Pi}$ that replace the above expectations with a transition risk mapping $\sigma : \mathcal{P}(\mathbb{R}) \times \mathcal{V} \to \mathbb{R}$ defined on the space $\mathcal{V}$ of all bounded and measurable functions from $\mathbb{R}$ to $\mathbb{R}$. We now specify the conditions for the mapping $\sigma(\cdot, \cdot)$ that will allow us use our earlier results.

Assumption 6.2. *The mapping σ is normalized: $\sigma(q, 0) = 0$, for all $q \in \mathcal{P}(\mathbb{R})$, and strongly monotonic with respect to stochastic dominance:*

$$\forall q^1, q^2 \in \mathcal{P}(\mathbb{R}),\ \forall v^1, v^2 \in \mathcal{V},\ (v^1; q^1) \preceq_{\text{st}} (v^2; q^2) \implies \sigma(q^1, v^1) \le \sigma(q^2, v^2).$$

Remark 6.3. *The general form derived in Theorem 5.6 is a family of risk mappings $\sigma_t : \mathbb{R} \times [0, 1] \times \mathcal{P}(\mathbb{R}) \times \mathcal{V} \to \mathbb{R}$, $t = 1, \dots, T$, which can model various degrees of risk-aversion for different periods t or different states (x, ξ). In our construction, we ignore the possibility of time- and state- dependence of risk aversion.*

Remark 6.4. *For a probability measure $q \in \mathcal{P}(\mathbb{R})$ that has f as the density function, with slight abuse of notation, we also write $\sigma(f, \cdot)$ instead of $\sigma(q, \cdot)$.*

Therefore, the risk measures $\{v_t^\pi\}_{t=1,\dots,T+1}$ are constructed in the following way:

$$\begin{gathered} v_{T+1}^\pi(h_{T+1}) = x_{T+1}; \\ v_t^\pi(h_t) = \begin{cases} x_t + \sigma\Big(\xi_t^\pi(h_t) f_1 + (1 - \xi_t^\pi(h_t)) f_2,\ x' \mapsto x' + v_{t+1}^\pi(h_t, x')\Big) & \text{if } \pi_t(h_t) = 0, \\ R + x_t + \sigma\Big(f_1,\ x' \mapsto x' + v_{t+1}^\pi(h_t, x')\Big) & \text{if } \pi_t(h_t) = 1, \end{cases} \\ t = T, \dots, 1. \end{gathered} \tag{45}$$

Here, the belief state $\xi_t^\pi(h_t)$ represents the conditional probability that $y_t = 1$ and is calculated recursively by (42) (and thus depends on the policy). Note that in the above formula, $\xi_t^\pi(h_t) f_1 + (1 - \xi_t^\pi(h_t)) f_2$ is the density function of the conditional distribution of x_{t+1} if the control is "continue".

Proposition 6.5. *for any Markov policy $\pi \in \Pi$, the functions v_t^π, $t = 1, \dots, T+1$ are memoryless, that is, $v_t^\pi(h_t) = v_t^\pi(x_t, \xi_t^\pi(h_t))$.*

Proof. By Theorem 5.6, the risk measures defined in (45) satisfy the Markov property of Definition 5.3, and thus the result follows from Corollary 5.7. □

6.3 Policy monotonicity

In this subsection, we assume that the transition risk mapping $\sigma : \mathcal{P}(\mathbb{R}) \times \mathcal{V} \to \mathbb{R}$ satisfies Assumption 6.2 and the following assumption.

Assumption 6.6. *σ is lower semicontinuous.*

The policy evaluation functions $\{v_t^\pi\}_{t=1,\dots,T+1}^{\pi\in\Pi}$ are defined recursively in (45). We want to study the properties of the optimal value functions (30). We obtain the following result.

Proposition 6.7. *The functions v_t^*, $t = 1, \dots, T+1$ are memoryless and satisfy the following dynamic programming equations:*

$$v_t^*(x,\xi) = x + \min\Big(R + \sigma(f_1, x' \mapsto v_{t+1}^*(x', 1-p));\\ \sigma(\xi f_1 + (1-\xi)f_2, x' \mapsto v_{t+1}^*(x', \Phi(\xi, x')))\Big),\\ x \in \mathbb{R}, \quad \xi \in [0,1], \quad t = 1, \dots, T, \quad (46)$$

with the final stage value $v_{T+1}^*(x,\xi) = x$. Moreover, an optimal Markov policy exists, which defined by the minimizers in the above dynamic programming equations.

Proof. We need to verify assumptions (i)-(iv) of Theorem 5.8. Assumption (i) is satisfied, because the density

$$q_t(x', y'|x, y, u) = \begin{cases} K^{[0]}_{y,y'} f_y(x') & \text{if } u = 0, \\ K^{[1]}_{y,y'} f_1(x') & \text{if } u = 1, \end{cases}$$

is uniformly bounded and continuous with respect to (x, u). Other assumptions are satisfied by construction. □

The following observation follows directly from (46).

Proposition 6.8. *The functions $v_t^*(x,\xi) - x$ depend only on ξ, and we can write*

$$v_t^*(x,\xi) = x + w_t^*(\xi), \quad t = 1, \dots, T+1.$$

The dynamic programming equations (46) *simplify as follows:*

$$w_t^*(\xi) = \min\Big\{R + \sigma(f_1, x' \mapsto x' + w_{t+1}^*(1-p));\\ \sigma(\xi f_1 + (1-\xi)f_2, x' \mapsto x' + w_{t+1}^*(\Phi(\xi, x')))\Big\},\\ x \in \mathbb{R}, \quad \xi \in [0,1], \quad t = 1, \dots, T, \quad (47)$$

with the final stage value $w_{T+1}^*(\,\cdot\,) = 0$.

We can now establish the monotonicity of w^*.

Theorem 6.9. *If $\frac{f1}{f2}$ is non-increasing, then the functions $w_t^* : [0,1] \to \mathbb{R}$, $t = 1, \dots, T+1$ are non-increasing.*

Proof. Clearly, w_{T+1}^* is non-increasing. Assume by induction that w_{t+1}^* is non-increasing. For any $\xi_1 \le \xi_2$ we have:

1. $\xi_1 f_1 + (1-\xi_1)f_2 \ge_{\text{st}} \xi_2 f_1 + (1-\xi_2)f_2$, because $f_1 \le_{\text{st}} f_2$.

2. For all x', we have $x' + w_{t+1}^*(\Phi(\xi_1, x')) \ge x' + w_{t+1}^*(\Phi(\xi_2, x'))$, as w_{t+1}^* is non-increasing and $\Phi(\cdot, x')$ is non-decreasing.

3. the mapping $x' \mapsto x' + w_{t+1}^*(\Phi(\xi, x'))$ is non-decreasing for all ξ. To show that, it is sufficient to establish that $x' \mapsto \Phi(\xi, x')$ is non-increasing, and this can be seen from the formula:

$$\frac{1}{\Phi(\xi, x')} = \frac{1}{p}\left(1 + \frac{f_2(x')}{f_1(x')}\left(\frac{1}{\xi} - 1\right)\right).$$

Thus

$$\begin{aligned}
&\sigma(\xi_1 f_1 + (1-\xi_1) f_2, x' \mapsto x' + w^*_{t+1}(\Phi(\xi_1, x'))) \\
&\quad \geq \sigma(\xi_1 f_1 + (1-\xi_1) f_2, x' \mapsto x' + w^*_{t+1}(\Phi(\xi_2, x'))) && \text{(because of 2.)} \\
&\quad \geq \sigma(\xi_2 f_1 + (1-\xi_2) f_2, x' \mapsto x' + w^*_{t+1}(\Phi(\xi_2, x'))) && \text{(because of 1. and 3.)}
\end{aligned}$$

which completes the induction step. □

The monotonicity assumption on $\frac{f_1}{f_2}$ is in fact a sufficient (but not necessary) condition for $f_1 \preceq_{\text{st}} f_2$. We illustrate this issue in the following examples.

Example 6.10 (Exponentially distributed costs)**.** We have

$$f_i(x) = \frac{1}{\theta_i} \exp\left(-\frac{x - m_i}{\theta_i}\right) \mathbb{1}_{\{x \geq m_i\}}, \quad i = 1, 2,$$

with $m_1, m_2, \theta_1, \theta_2 \geq 0$. To have $f_1 \preceq_{\text{st}} f_2$, we must have $m_1 \leq m_2$ and $\theta_1 \geq \theta_2$. Thus the non-increasingness of $\frac{f_1}{f_2}$ is equivalent to $f_1 \preceq_{\text{st}} f_2$, as

$$\frac{f_1(x)}{f_2(x)} = \begin{cases} +\infty, & \text{for } m_1 \leq x < m_2, \text{ if } m_1 < m_2; \\ 0, & \text{for } m_2 \leq x < m_1, \text{ if } m_1 > m_2; \\ \dfrac{\theta_2}{\theta_1} \exp\left(-\dfrac{x - m_1}{\theta_1} + \dfrac{x - m_2}{\theta_2}\right), & \text{for } x \geq \max(m_1, m_2). \end{cases}$$

Example 6.11 (Normally distributed costs)**.** *Consider the normal distributions truncated to* $[0, +\infty)$

$$f_i(x) = \frac{\exp\left(-\frac{1}{2}\left(\frac{x-\mu_i}{\sigma_i}\right)^2\right)}{\int_0^\infty \exp\left(-\frac{1}{2}\left(\frac{z-\mu_i}{\sigma_i}\right)^2\right) dz}, \quad i = 1, 2,$$

where $\sigma_1, \sigma_2 > 0$ *and* $\mu_1, \mu_2 \in \mathbb{R}$*. We have*

$$\frac{f_1(x)}{f_2(x)} = \textit{const.} \cdot \exp\left(-\frac{1}{2}\left(\frac{x-\mu_1}{\sigma_1}\right)^2 + \frac{1}{2}\left(\frac{x-\mu_2}{\sigma_2}\right)^2\right),$$

so the derivative of $\ln \frac{f_1}{f_2}$ *is* $-\frac{x-\mu_1}{\sigma_1^2} + \frac{x-\mu_2}{\sigma_2^2}$*. A necessary and sufficient condition for* $\frac{f_1}{f_2}$ *being non-increasing on* $[0, +\infty)$ *is*

$$\sigma_1 \leq \sigma_2 \ \textit{ and } \ \frac{\mu_1}{\sigma_1^2} \leq \frac{\mu_2}{\sigma_2^2}.$$

Example 6.12 (Uniformly distributed costs)**.** *We set*

$$f_i = \frac{1}{M_i - m_i} \mathbb{1}_{[m_i, M_i]}, \quad i = 1, 2.$$

To have $f_1 \preceq_{\text{st}} f_2$*, we must have* $m_1 \leq m_2$ *and* $M_1 \leq M_2$*, and then*

$$\frac{f_1(x)}{f_2(x)} = \begin{cases} +\infty, & \textit{for } m_1 \leq x \leq m_2; \\ \dfrac{M_2 - m_2}{M_1 - m_1}, & \textit{for } m_2 < x \leq M_1; \\ 0, & \textit{for } x \geq M_2, \end{cases}$$

is non-increasing.

As a consequence of Theorem 6.9, we have the following threshold property of the policy.

Theorem 6.13. *Under the assumptions of Theorem 6.9, there exist thresholds $\xi_t^* \in [0,1]$, $t = 1,\dots,T$ such that the policy*

$$u_t^* = \begin{cases} 0 & \text{if } \xi_t > \xi_t^*, \\ 1 & \text{if } \xi_t \le \xi_t^*, \end{cases}$$

is optimal.

Proof. Suppose ξ is such that replacement at time t is optimal:

$$R + \sigma(f_1, x' \mapsto x' + w_{t+1}^*(1-p)) \le \sigma(\xi f_1 + (1-\xi) f_2, x' \mapsto x' + w_{t+1}^*(\Phi(\xi, x'))).$$

Then for any $\zeta \le \xi$, we have $\xi f_1 + (1-\xi) f_2 \preceq_{\mathrm{st}} \zeta f_1 + (1-\zeta) f_2$ and $\Phi(\xi, x') \ge \Phi(\zeta, x')$. Consequently,

$$\begin{aligned} R &+ \sigma(f_1, x' \mapsto x' + w_{t+1}^*(1-p)) \\ &\le \sigma(\xi f_1 + (1-\xi) f_2, x' \mapsto x' + w_{t+1}^*(\Phi(\xi, x'))) \\ &\le \sigma(\zeta f_1 + (1-\zeta) f_2, x' \mapsto x' + w_{t+1}^*(\Phi(\xi, x'))) \\ &\le \sigma(\zeta f_1 + (1-\zeta) f_2, x' \mapsto x' + w_{t+1}^*(\Phi(\zeta, x'))), \end{aligned}$$

and replacement is optimal for ζ as well. □

6.4 Numerical illustration

In this section, we solve the problem in the special case where f_1 and f_2 are density functions of $\mathbb{U}(m_1, M_1)$ and $\mathbb{U}(m_2, M_2)$ with $m_1 \le m_2 \le M_1 \le M_2$, as in Example 6.12. Then the Bayes operator is piece-wise constant with respect to x':

$$\Phi(\xi, x') = \begin{cases} 1-p, & \text{if } m_1 \le x' \le m_2; \\ \dfrac{(1-p)\xi(M_2 - m_2)}{\xi(M_2 - m_2) + (1-\xi)(M_1 - m_1)} := \hat{\phi}(\xi), & \text{if } m_2 \le x' \le M_1; \\ 0, & \text{if } M_1 \le x' \le M_2. \end{cases}$$

The conditional distribution of x' given ξ is described by the density function $\xi f_1 + (1-\xi) f_2$, which is also constant in each of the three intervals $[m_1, m_2)$, $[m_2, M_1]$ and $(M_1, M_2]$, with the probabilities amassed in each of the three intervals

$$\begin{aligned} q_1(\xi) &= \frac{\xi(m_2 - m_1)}{M_1 - m_1}, \\ q_2(\xi) &= (M_1 - m_2)\left(\frac{\xi}{M_1 - m_1} + \frac{1-\xi}{M_2 - m_2}\right), \\ q_3(\xi) &= \frac{(1-\xi)(M_2 - M_1)}{M_2 - m_2}. \end{aligned}$$

We use mean-semi-deviation for σ, i.e,

$$\sigma(q, v) = \mathbb{E}_q(v) + \gamma \mathbb{E}_q\left(v - \mathbb{E}_q(v)\right)^+,$$

which is strongly monotonic with respect to stochastic order and lower semi-continuous with respect to (q, v). Then the dynamic programming equations (47) for $t = 1, \ldots, T$ become:

$$w_t^*(\xi) = \min\Big\{R + E_t^*(1) + \mathbb{E}_{f_1}(x' \mapsto x' + w_{t+1}^*(1-p) - E_t^*(1))^+; \\ E_t^*(\xi) + \mathbb{E}_{\xi f_1 + (1-\xi) f_2}(x' \mapsto x' + w_{t+1}^*(\Phi(\xi, x')) - E_t^*(\xi))^+, \Big\}, \quad (48)$$

where

$$\begin{aligned} E_t^*(\xi) &:= \mathbb{E}_{\xi f_1 + (1-\xi) f_2}(x' \mapsto x' + w_{t+1}^*(\Phi(\xi, x'))) \\ &= q_1(\xi)\left(\frac{m_1 + m_2}{2} + w_{t+1}^*(1-p)\right) \\ &\quad + q_2(\xi)\left(\frac{m_2 + M_1}{2} + w_{t+1}^*(\hat{\phi}(\xi))\right) \\ &\quad + q_3(\xi)\left(\frac{M_1 + M_2}{2} + w_{t+1}^*(0)\right). \end{aligned} \quad (49)$$

As $x' \mapsto x' + w_{t+1}^*(\Phi(\xi, x')) - E_t^*(\xi)$ is linear in each of the intervals $[m_1, m_2]$, $[m_2, M_1]$ and $[M_1, M_2]$, we have

$$\begin{aligned} &\mathbb{E}_{\xi f_1 + (1-\xi) f_2}(x' \mapsto x' + w_{t+1}^*(\Phi(\xi, x')) - E_t^*(\xi))^+ \\ &\quad = q_1(\xi)\,\theta(m_1, m_2, E_t^*(\xi) - w_{t+1}^*(1-p)) \\ &\quad + q_2(\xi)\,\theta(m_2, M_1, E_t^*(\xi) - w_{t+1}^*(\hat{\phi}(\xi))) \\ &\quad + q_3(\xi)\,\theta(M_1, M_2, E_t^*(\xi) - w_{t+1}^*(0)), \end{aligned} \quad (50)$$

where, for $a_1 \le a_2$,

$$\theta(a_1, a_2, a_3) := \frac{\int_{a_1}^{a_2} (\cdot - a_3)^+}{a_2 - a_1} = \begin{cases} \frac{1}{2}(a_1 + a_2), & \text{if } a_3 <= a_1; \\ \frac{1}{2}(a_3 + a_2)\frac{a_2 - a_3}{a_2 - a_1}, & \text{if } a_1 < a_3 < a_2; \\ 0, & \text{if } a_3 >= a_2. \end{cases}$$

For any t and any ξ, if we know $w_{t+1}^*(1-p)$, $w_{t+1}^*(\hat{\phi}(\xi))$ and $w_{t+1}^*(0)$, then the computation of $w_t^*(\xi)$ can be accomplished in three steps:

1. Compute $E_t^*(1)$ and $E_t^*(\xi)$ by (49).

2. Compute $\mathbb{E}_{f_1}(x' \mapsto x' + w_{t+1}^*(1-p) - E_t^*(1))^+$ and $\mathbb{E}_{\xi f_1 + (1-\xi) f_2}(x' \mapsto x' + w_{t+1}^*(\Phi(\xi, x')) - E_t^*(\xi))^+$ by (50).

3. Compute $w_t^*(\xi)$ using the dynamic programming equation (48).

Since we have the final stage value $w_{T+1}^* = 0$, all $w_t^*(\xi)$ can be easily calculated by recursion backward in time.

In Figure 1 we present the value functions and optimal policies for an example with $m_1 = 0$, $m_2 = 80$, $M_1 = 100$, $M_2 = 500$, $p = 0.2$, $T = 6$, $R = 50$, and $\gamma = 0.9$. In Figure 2 we display the distribution of the total cost obtained by simulating 100,000 runs of the system with both policies.

We see that the application of the risk-averse model increases the threshold values ξ_t^* of the optimal policies and results in a significantly less dispersed distribution of the total cost.

References

[1] A. Arlotto, N. Gans, and J. M. Steele. Markov decision problems where means bound variances. *Operations Research*, 62(4):864–875, 2014.

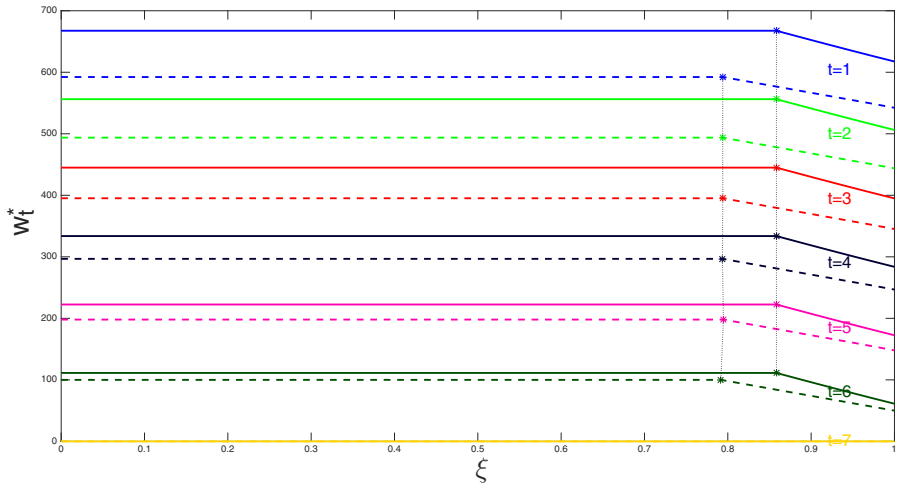

Figure 1: The value function for the risk-neutral model (dashed) and the risk-averse model (solid). The stars denote the critical value ξ^* below which replacement is optimal.

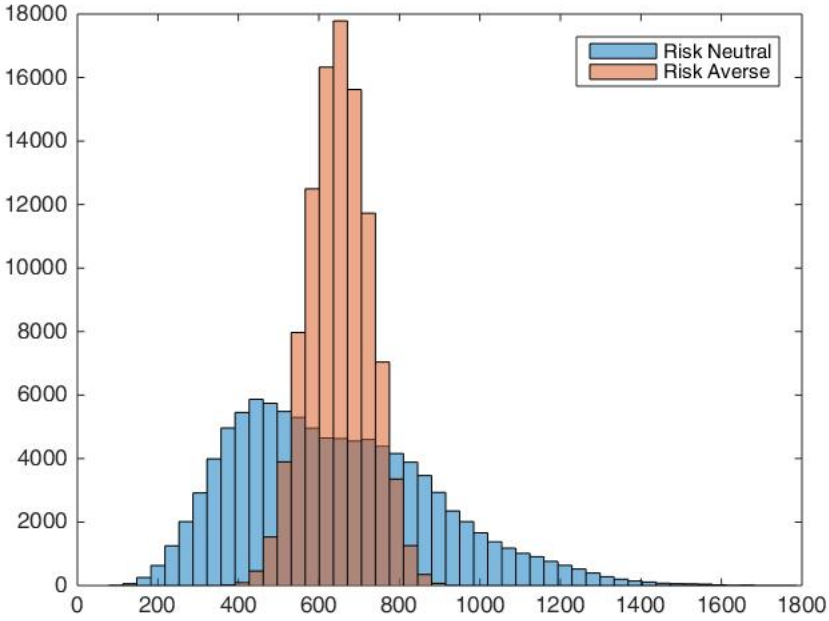

Figure 2: Empirical distribution of the total cost for the risk-neutral model (blue) and the risk-averse model (orange).

[2] P. Artzner, F. Delbaen, J.-M. Eber, D. Heath, and H. Ku. Coherent multiperiod risk adjusted values and Bellman's principle. *Annals of Operations Research*, 152:5–22, 2007.

[3] J.-P. Aubin and H. Frankowska. *Set-valued analysis*. Birkhäuser, Boston, MA, 2009.

[4] N. Bäuerle and U. Rieder. *Markov decision processes with applications to finance*. Universitext. Springer, Heidelberg, 2011.

[5] N. Bäuerle and U. Rieder. More risk-sensitive Markov decision processes. *Mathematics of Operations Research*, 39(1):105–120, 2013.

[6] D. P. Bertsekas and S. E. Shreve. *Stochastic optimal control*, volume 139 of *Mathematics in Science and Engineering*. Academic Press, New York-London, 1978.

[7] Ö. Çavus and A. Ruszczyński. Computational methods for risk-averse undiscounted transient Markov models. *Operations Research*, 62(2):401–417, 2014.

[8] Ö. Çavus and A. Ruszczyński. Risk-averse control of undiscounted transient Markov models. *SIAM Journal on Control and Optimization*, 52(6):3935Ű–3966, 2014.

[9] Z. Chen, G. Li, and Y. Zhao. Time-consistent investment policies in Markovian markets: a case of mean-variance analysis. *J. Econom. Dynam. Control*, 40:293–316, 2014.

[10] P. Cheridito, F. Delbaen, and M. Kupper. Dynamic monetary risk measures for bounded discrete-time processes. *Electronic Journal of Probability*, 11:57–106, 2006.

[11] P. Cheridito and M. Kupper. Composition of time-consistent dynamic monetary risk measures in discrete time. *International Journal of Theoretical and Applied Finance*, 14(01):137–162, 2011.

[12] F. Coquet, Y. Hu, J. Mémin, and S. Peng. Filtration-consistent nonlinear expectations and related g-expectations. *Probability Theory and Related Fields*, 123(1):1–27, 2002.

[13] S. P. Coraluppi and S. I. Marcus. Risk-sensitive and minimax control of discrete-time, finite-state Markov decision processes. *Automatica*, 35(2):301–309, 1999.

[14] P. Dai Pra, L. Meneghini, and W. J. Runggaldier. Explicit solutions for multivariate, discrete-time control problems under uncertainty. *Systems & Control Letters*, 34(4):169–176, 1998.

[15] E. V. Denardo and U. G. Rothblum. Optimal stopping, exponential utility, and linear programming. *Math. Programming*, 16(2):228–244, 1979.

[16] G. B. Di Masi and Ł. Stettner. Risk-sensitive control of discrete-time Markov processes with infinite horizon. *SIAM J. Control Optim.*, 38(1):61–78, 1999.

[17] E. A. Feinberg, P. O. Kasyanov, and M. Z. Zgurovsky. Partially observable total-cost Markov decision processes with weakly continuous transition probabilities. preprint, arXiv:1401.2168v1, 2014.

[18] J. A. Filar, L. C. M. Kallenberg, and H.-M. Lee. Variance-penalized Markov decision processes. *Math. Oper. Res.*, 14(1):147–161, 1989.

[19] H. Föllmer and I. Penner. Convex risk measures and the dynamics of their penalty functions. *Statistics & Decisions*, 24(1/2006):61–96, 2006.

[20] K. Hinderer. *Foundations of non-stationary dynamic programming with discrete time parameter*. Lecture Notes in Operations Research and Mathematical Systems, Vol. 33. Springer-Verlag, Berlin-New York, 1970.

[21] R. A. Howard and J. E. Matheson. Risk-sensitive Markov decision processes. *Management Sci.*, 18:356–369, 1971/72.

[22] S. C. Jaquette. Markov decision processes with a new optimality criterion: discrete time. *Ann. Statist.*, 1:496–505, 1973.

[23] S. C. Jaquette. A utility criterion for Markov decision processes. *Management Sci.*, 23(1):43–49, 1975/76.

[24] A. Jobert and L. C. G. Rogers. Valuations and dynamic convex risk measures. *Mathematical Finance*, 18(1):1–22, 2008.

[25] S. Klöppel and M. Schweizer. Dynamic indifference valuation via convex risk measures. *Math. Finance*, 17(4):599–627, 2007.

[26] M. Kupper and W. Schachermayer. Representation results for law invariant time consistent functions. *Mathematics and Financial Economics*, 2(3):189–210, 2009.

[27] K. Kuratowski and C. Ryll-Nardzewski. A general theorem on selectors. *Bull. Acad. Polon. Sci. Ser. Sci. Math. Astronom. Phys*, 13(1):397–403, 1965.

[28] S. Levitt and A. Ben-Israel. On modeling risk in Markov decision processes. In *Optimization and related topics (Ballarat/Melbourne, 1999)*, volume 47 of *Appl. Optim.*, pages 27–40. Kluwer Acad. Publ., Dordrecht, 2001.

[29] K. Lin and S. I. Marcus. Dynamic programming with non-convex risk-sensitive measures. In *American Control Conference (ACC), 2013*, pages 6778–6783. IEEE, 2013.

[30] S. Mannor and J. N. Tsitsiklis. Algorithmic aspects of mean-variance optimization in Markov decision processes. *European J. Oper. Res.*, 231(3):645–653, 2013.

[31] S. I. Marcus, E. Fernández-Gaucherand, D. Hernández-Hernández, S. Coraluppi, and P. Fard. Risk sensitive Markov decision processes. In *Systems and control in the twenty-first century (St. Louis, MO, 1996)*, volume 22 of *Progr. Systems Control Theory*, pages 263–279. Birkhäuser, Boston, MA, 1997.

[32] W. Ogryczak and A. Ruszczyński. From stochastic dominance to mean-risk models: Semideviations as risk measures. *European Journal of Operational Research*, 116(1):33–50, 1999.

[33] W. Ogryczak and A. Ruszczyński. On consistency of stochastic dominance and mean–semideviation models. *Mathematical Programming*, 89(2):217–232, 2001.

[34] G.Ch. Pflug and W. Römisch. *Modeling, Measuring and Managing Risk*. World Scientific, Singapore, 2007.

[35] Z. Porosiński, K. Szajowski, and S. Trybuła. Bayes control for a multidimensional stochastic system. *System Sciences*, 11:51–64, 1985.

[36] F. Riedel. Dynamic coherent risk measures. *Stochastic Processes and Their Applications*, 112:185–200, 2004.

[37] R T. Rockafellar and R. J.-B. Wets. *Variational analysis*, volume 317. Springer, Berlin, 1998.

[38] B. Roorda, J. M. Schumacher, and J. Engwerda. Coherent acceptability measures in multiperiod models. *Mathematical Finance*, 15(4):589–612, 2005.

[39] W. J. Runggaldier. Concepts and methods for discrete and continuous time control under uncertainty. *Insurance: Mathematics and Economics*, 22(1):25–39, 1998.

[40] A. Ruszczyński. Risk-averse dynamic programming for Markov decision processes. *Math. Program.*, 125(2, Ser. B):235–261, 2010.

[41] A. Ruszczyński and A. Shapiro. Conditional risk mappings. *Mathematics of Operations Research*, 31:544–561, 2006.

[42] G. Scandolo. *Risk Measures in a Dynamic Setting*. PhD thesis, Università degli Studi di Milano, Milan, Italy, 2003.

[43] A. Shapiro. Time consistency of dynamic risk measures. *Operations Research Letters*, 40(6):436–439, 2012.

[44] Y. Shen, W. Stannat, and K. Obermayer. Risk-sensitive Markov control processes. *SIAM Journal on Control and Optimization*, 51(5):3652–3672, 2013.

[45] S. Weber. Distribution-invariant risk measures, information, and dynamic consistency. *Mathematical Finance*, 16(2):419–441, 2006.

[46] D. J. White. Mean, variance, and probabilistic criteria in finite Markov decision processes: a review. *J. Optim. Theory Appl.*, 56(1):1–29, 1988.